Foundational Causality

John A Challoner, Aug., 2026

Abstract

Causality is fundamental to systems theory because systems exist and operate through interactions that produce, maintain and alter states, processes and organisation. This paper develops a foundational account of causality from a Critical Realist and physicalist perspective, building upon previously established concepts of entities, states, processes, transfers and systems. Causal influence is proposed to be physically instantiated through transfers of matter, energy or physically embodied information between processes. These interactions concatenate into causal networks whose organisation may be transient or persistent and may include necessary, sufficient, enabling and inhibiting causal contributions.

The underlying causal organisation is represented as an alternating sequence of processes and transfers. Two complementary epistemic perspectives upon this sequence are developed: Process–Transfer–Process (PTP), which foregrounds causal interaction, and Transfer–Process–Transfer (TPT), which foregrounds system functioning through the transformation of inputs into outputs. Their duality provides a conceptual bridge between causal theory and systems theory. Persistent causal organisations establish continuing possibilities for interaction and may exhibit emergent properties, behaviours and capabilities through which they can be recognised as systems.

The paper further examines causal leverage, whereby relatively small inputs may enable, inhibit or redirect much larger independently available causal flows, providing a physical basis for control and agency. Finally, it distinguishes causal reality from causal knowledge. Recurring causal relationships and organisations may be recognised and represented epistemically as hypotheses, explanations and theories, whose explanatory power depends partly upon the recurrence of the causal organisation they represent. The resulting framework treats systems and causality as complementary aspects of the same underlying reality and provides foundations for subsequent work on emergence, cognition, agency, hypothesis testing and systems investigation.

1. Introduction

Causality is fundamental to systems theory because systems do not exist in isolation. Their components interact, systems interact with other systems and with their environments, and these interactions can change states, behaviour and organisation. Understanding systems therefore requires an account not only of what systems are, but also of how one part of reality can influence another.

The previous paper in this series, Foundations of Systems Theory from a Critical Realist Perspective (Challoner, 2026), established a set of foundational concepts including entities, properties, states, change, relationships, interactions, processes, transfers and systems. This paper builds upon those foundations by examining the causal relationships through which change occurs and through which systems interact and operate.

This treatment is consistent with the Critical Realist view that causal mechanisms and powers belong to reality rather than being reducible to observed regularities between events (Bhaskar, 1978). The central position adopted here is that causality is physically instantiated through processes and transfers. Where one process influences another, a physical causal pathway must connect them. What is transferred may conveniently be described as matter, energy or information, although information is itself physically embodied as structure in matter or energy (Landauer, 1991).

For example, pressing a doorbell button changes the configuration of an electrical circuit, allowing electrical energy to be transferred to the doorbell, where further processes produce sound. Describing the button press as causing the sound identifies a causal relationship; identifying the intervening processes and transfers helps explain how that relationship is physically realised.

Causal relationships concatenate into larger causal networks. Some are transient, while others persist and form part of the circumstances, conditions and constraints within which further causation occurs. Persistent configurations establish continuing possibilities for interaction, while processes and transfers realise particular causal pathways. Relatively coherent causal organisations may exhibit emergent properties, behaviours and outputs and thereby be recognised as systems.

The paper also distinguishes causal reality from its epistemic representation. The underlying causal organisation can be represented as an alternating sequence of processes and transfers, while PTP provides an event-based perspective upon causal interactions and TPT a systems-based perspective upon process functioning. These complementary representations provide a conceptual bridge between causal theory and systems theory. Causal relationships and recurring causal organisations may subsequently be recognised and represented epistemically as hypotheses, explanations and theories, whose testing and refinement can provide progressively more reliable causal knowledge.

Finally, the paper considers causal leverage, through which relatively small causal inputs may enable, inhibit, redirect or otherwise control much larger independently available causal flows. Causal leverage arises from the position and effect of an input within a causal organisation rather than simply from the magnitude of the initiating transfer. This provides a physical connection between foundational causality and the later treatment of agency, through which relatively small actions within an agent may ultimately contribute to substantial changes in the external world.

The argument developed here therefore treats systems theory and causal theory as strongly interrelated. Causal organisation helps explain how systems form, persist, interact and produce effects, while systems provide persistent organisations through which causality is realised. This common causal foundation provides a basis for subsequent treatments of emergence, cognition, agency, hypothesis testing and systems diagnosis.

2. Causality and Change

2.1 Cause and Effect

Causality is commonly described as a relationship in which a cause produces an effect. However, causes and effects should not be regarded as particular kinds of concrete entity. Cause and effect are epistemic roles attributed to causal phenomena when portions of causal reality are represented for explanatory purposes.

At this stage these roles will be used in their familiar, relatively compressed form. Later sections will decompose causal relationships into their underlying processes and transfers and show more precisely what may occupy the roles of cause and effect.

The same causal phenomenon may be represented as an effect in relation to preceding causation and as a cause in relation to subsequent causation. For example, increased river flow may be represented as an effect of rainfall and subsequently as a cause contributing to downstream flooding.

This means that causal relationships rarely exist as isolated pairs. They normally form parts of longer causal sequences in which effects become causes of subsequent effects:

Cause → Effect/Cause → Effect/Cause → Effect

The terms cause and effect therefore identify roles attributed to causal phenomena within the particular causal relationship being represented. Changing the boundary or granularity of analysis may change which role is attributed to them.

This also means that causality is better understood as a continuing network of interactions rather than as a collection of independent cause-and-effect pairs. Individual causal relationships are useful epistemic representations of portions of that larger causal reality. This provides the basis for examining causal chains and networks later in the paper.

2.2 State Change

Causality and change are closely related. A causal interaction occurs when an influence from one entity or process contributes to a change in another. In systems terms, this can be understood as a transfer contributing to a change in the state of the receiving entity or process.

Change should not, however, be understood only as the alteration of an existing property. What is represented as a causal effect may take several forms. It may:

- bring an entity or process into existence;
- bring an entity or process to an end;
- alter one or more of its properties and therefore its state; or
- contribute to maintaining its existing state or continued operation.

The last case is particularly important. Causation can occur even where there is no obvious change at the level being observed. For example, a pump may continuously replace water leaving a tank and thereby maintain a constant water level. The apparently unchanged state of the tank does not indicate an absence of causation. It results from continuing causal processes whose effects balance one another.

What is represented as an effect should therefore not be equated simply with observable change. It may instead be the production, alteration, termination or maintenance of a state or process. What appears macroscopically to be stability may consequently depend upon continuing causal activity at a finer level.

This connects causality directly with the concepts of state and change developed in the previous paper in this series. Causality provides the means by which states are produced and altered, while continuing causal processes can also explain how particular states are maintained.

2.3 Causes as Effects and Effects as Causes

Because cause and effect are epistemic roles attributed within causal relationships, what is represented as an effect in one relationship may participate in subsequent causation and subsequently be represented as a cause. Causal influence can thereby propagate through reality.

Causation should therefore not be understood simply as a collection of separate relationships of the form **A causes B**. Individual causal relationships are portions of larger and potentially continuing causal sequences. Where these sequences branch, converge or interact, they form the causal networks considered in Section 5.

This continuity is important because the cause selected to explain a particular effect may itself have preceding causes, while the effect selected as the endpoint of an explanation may have further consequences. Cause and effect therefore depend partly upon which portion of a continuing causal sequence is being considered.

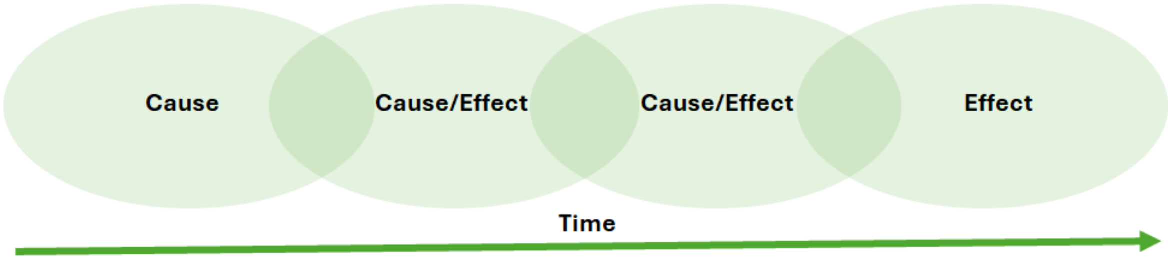

Figure 1. *From Cause and Effect to a Causal Sequence*

Cause and effect therefore depend partly upon which portion of continuing causal reality is selected for representation and upon the granularity at which it is represented. Later sections will decompose this causal continuity into processes and transfers and show how PTP events and the functioning or non-functioning of processes represented through TPT may occupy these causal roles.

3. The Physical Basis of Causality

3.1 Causal Interaction

Conventional causal descriptions tend to identify causes and effects with observable events, processes, states or changes. This may partly reflect their relative epistemic accessibility. Such phenomena are often comparatively easy to observe and distinguish, whereas the transfers through which causal influence propagates may be less apparent or require further investigation to identify. Conventional cause-and-effect descriptions may therefore compress an underlying causal organisation by identifying its more readily observable elements while leaving the transfers connecting them implicit.

Events that occur together, or one after another, are not necessarily causally related. Their association may be coincidental, or both may result from another cause. Temporal succession and statistical association can therefore provide evidence from which a causal relationship may be suspected, but do not by themselves demonstrate that one thing has influenced another (Pearl, 2009).

A genuine causal relationship requires physical interaction between the relevant portions of causal reality. Within the framework developed here, causal influence between processes is physically instantiated through a transfer.

For example, switching on a lamp is normally followed by the lamp illuminating. The temporal sequence alone does not establish causation. The causal relationship is demonstrated more fully when the intervening interaction is identified: operating the switch changes the circuit, allowing electrical energy to be transferred to the lamp, where further processes produce light.

This distinction becomes especially important in complex systems. Two variables may repeatedly change together without either causing the other. For example, hot weather may be associated with both increased ice-cream sales and increased use of air conditioning. The association between the latter two does not imply a causal interaction between them; both are influenced by another causal factor.

Causal analysis therefore seeks more than recurring association. It seeks to identify the processes and transfers through which causal influence propagates. This provides the physical connection between cause and effect and forms the basis for the more detailed treatment of causal transfer that follows.

3.2 Causal Transfer

For one process to influence another, something must pass between them. This is referred to here as a causal transfer. An output from one process becomes available as an input to another, providing the physical connection through which causal influence propagates.

Causal transfer may involve matter, energy or information. Water flowing through a pipe transfers matter and energy; a collision transfers energy between moving objects; light transfers electromagnetic energy; and a spoken instruction transfers information embodied in sound waves. Although these examples differ greatly, each involves something physically passing from one process to another.

This physical treatment has important antecedents in causal-process theories, particularly the work of Salmon and Dowe, which sought to understand causation through physical processes, interactions and the transmission or possession of conserved quantities (Dowe, 2000; Salmon, 1984).

This leads to a foundational principle:

No causal influence can pass between separate processes without a physical transfer connecting them.

The transfer need not itself provide all of the matter, energy or information responsible for the resulting effect. A small transfer may trigger, enable, inhibit or redirect much larger flows already available within the receiving system. Pressing an electrical switch, for example, requires very little energy but can control a much larger flow of electrical energy supplied from elsewhere. This distinction will become important when causal leverage is considered later in the paper.

Causal transfer also explains why causation takes time. Matter and energy cannot propagate instantaneously, and processes receiving them also require time to respond. Causal relationships therefore possess both spatial and temporal structure.

The familiar cause-and-effect representation introduced in Section 2 can therefore now begin to be decompressed. At its most elementary, the underlying physical causality can be represented as:

Process → Transfer → Process

Extended causation arises when this pattern continues, producing the alternating sequence of processes and transfers that forms larger causal chains and networks.

3.3 Matter, Energy and Information

Causal transfers may conveniently be described as transfers of matter, energy or information. This distinction is useful because it identifies what is significant about a transfer, but it should not be taken to imply that these are three separate fundamental substances.

Matter may pass directly between processes. Water entering a plant through its roots, food entering an organism, and components moving along a production line are examples. Energy may likewise be transferred, as occurs when one object collides with another, electrical energy passes through a circuit, or electromagnetic radiation reaches a receiving surface.

Information requires a slightly different treatment. Within the physicalist framework adopted here, information is not an abstract substance that can travel independently of the physical world. Information is structure embodied in matter or energy (Landauer, 1991). An informational transfer therefore requires a physical carrier whose structure can influence processes within the receiving system.

A spoken sentence provides a simple example. Sound waves transfer energy through the air, but their particular pattern also carries information. The energy involved may be very small, yet the structure of that energy can produce substantial effects when it interacts with the processes of a person who understands the sentence. Similarly, information may be embodied in electrical signals, chemical structures, printed marks or patterns of light.

The distinction between matter, energy and information therefore describes different aspects of physical causal transfer. Matter and energy provide the physical substrate, while information concerns the organisation of that substrate. Informational causation consequently remains physically instantiated: information can have causal effects because the matter or energy in which it is embodied interacts with processes capable of responding to that organisation.

3.4 Source, Transfer and Recipient

A causal transfer involves three related elements: a source, something transferred, and a recipient. The resulting effect depends upon all three.

First, whatever is transferred must be physically available to the source. Matter cannot be transferred unless the source contains or receives that matter, and energy cannot be transferred without an available energy source. Similarly, informational structure cannot be transmitted unless the relevant structure is embodied within the causal organisation producing the transfer.

Second, the transfer must provide a physical connection between source and recipient. What reaches the recipient need not be identical to what originally existed within the source because it may be transformed during transmission. Nevertheless, there must be a continuous causal pathway connecting them.

Third, the recipient must possess processes capable of responding to the transfer. The same input may therefore produce different effects in different recipients. Sunlight falling on a solar panel may generate electricity, while the same sunlight falling on a stone produces mainly heating. The difference lies not simply in the incoming energy but in the organisation and processes of the recipient.

This is especially clear for informational transfers. A spoken instruction can influence someone who hears and understands it, but may have little or no informational effect on someone who does not understand the language. The physical transfer of sound occurs in both cases; the causal consequences differ because the receiving processes differ.

An effect is therefore not determined by the transfer alone. It arises from the interaction between what the source provides, what is transferred, and how the recipient responds. Causality consequently depends upon the organisation of both the source and the recipient as well as the causal pathway connecting them.

3.5 Transformation and Integration

A causal transfer does not necessarily pass unchanged through the receiving process. Once received, matter, energy or information may be transformed, decomposed, reorganised or incorporated into the recipient. The resulting effect therefore depends not only upon what is transferred, but also upon what the receiving process does with it.

Transformation occurs when the received input is converted into a different form. For example, a solar panel converts energy carried by sunlight into electrical energy. Similarly, a plant uses photosynthesis to transform incoming energy and matter into chemical structures that contribute to its growth.

Decomposition and reorganisation are particularly important where matter is transferred. Food entering an organism is not simply added to it unchanged. Digestion breaks complex materials into simpler components, some of which are subsequently reorganised and incorporated into the organism. The transferred material thereby becomes part of a new causal organisation.

Information may also be transformed. A message received by a person is interpreted through existing cognitive processes and may be combined with previous knowledge before influencing subsequent behaviour. The resulting effect is therefore not simply a copy of the informational structure received.

In some cases, transferred components become incorporated relatively permanently into the recipient. This is particularly important in assembly, where existing components are causally brought together to form a new organisation. Assembly will therefore be considered later as a distinctive form of causal integration.

Causal effects should consequently not be understood as simple transmission from source to recipient. The source provides a transfer, but the receiving process helps determine the resulting effect through the ways in which that transfer is transformed and integrated.

3.6 Causal Delay and Spatio-temporal Continuity

Causal influence cannot propagate instantaneously. Physical transfers of matter, energy or information require time to travel between processes, and the receiving process also requires time to respond. Every causal relationship therefore involves some causal delay.

The delay may be extremely short or very long. Light from the Sun takes about eight minutes to reach the Earth, while changes within an ecological or social system may take years before their consequences become apparent. The timescale differs, but the underlying principle is the same: causes must precede the effects they produce.

Causality also requires spatial continuity. A cause does not need to be immediately adjacent to its effect, but there must be a continuous causal pathway connecting them. A radio transmitter, for example, can influence a receiver many kilometres away because electromagnetic radiation propagates through the intervening space. Distance separates the two systems, but the physical transfer connects them.

Processes and transfers therefore give causal networks an inherent spatio-temporal order. A transfer originates through one process, propagates through space and time, and interacts with another process. That process may subsequently produce further transfers, extending the causal sequence.

This provides an important constraint upon causal explanation. If no physically possible pathway can connect a proposed cause to an effect within the available time, the proposed causal relationship must be questioned. Conversely, identifying the intervening processes, transfers and delays helps establish how causal influence can propagate through a network.

Causal networks are therefore not merely logical arrangements of causes and effects. They are physically connected organisations extending through space and time.

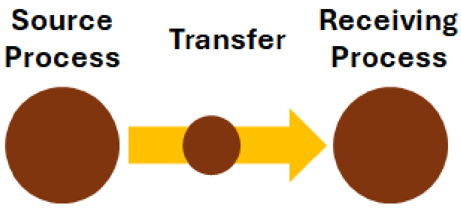

Figure 2. Physical Causal Transfer Between Processes.
A causal interaction in which matter, energy or physically embodied information is transferred from one process to another. The arrow represents the transfer and the smaller circle represents that which is transferred.

4. Necessary, Sufficient, Enabling and Inhibiting Causation

4.1 Foundational Principles

Necessity and sufficiency are always relative to a specified process or effect and its inputs or causal transfers. If a particular transfer is required for the process to operate, the transfer is necessary irrespective of where it originates. Where the complete set of necessary transfers is present, and no relevant inhibitor prevents the process, their combination is sufficient for the process to occur and produce the specified effect.

4.2 Necessary Causes

A necessary cause is a causal contribution without which a specified process or effect cannot occur. Most processes require several such contributions. Each is necessary, while the complete set of necessary inputs, in the absence of effective inhibition, is sufficient for the process to occur.

That contribution is physically instantiated through transfers of matter, energy or information. Where the necessary cause is described as a condition, the condition may itself be an ongoing causal network that provides a necessary causal contribution to the process producing the effect.

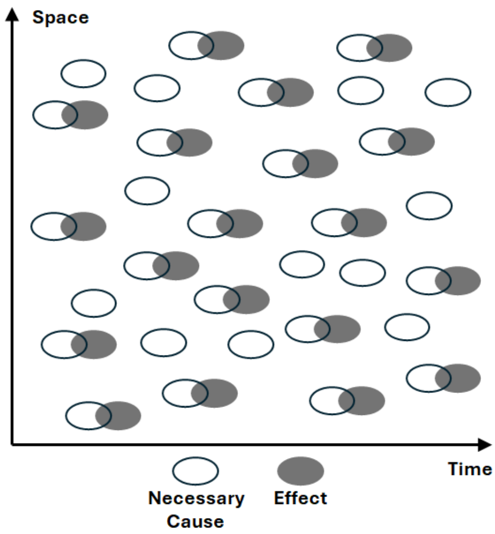

Figure 3. Necessary Cause.
Conventionally, a necessary cause is one in the absence of which the specified effect cannot occur (Mackie, 1974).

For example, oxygen is necessary for an ordinary combustion process. In its absence, combustion cannot occur. However, the presence of oxygen does not itself cause combustion: suitable fuel, an ignition source and other appropriate conditions are also required.

Identifying necessary causes is particularly useful in causal analysis because their absence can explain why an expected effect has not occurred. Conversely, removing a necessary cause can provide a means of preventing an unwanted effect.

4.3 Sufficient Causes

A sufficient cause is a cause that, when present, always produces the stated effect. A sufficient cause therefore differs from a necessary cause. A necessary cause is required for an effect but may not be enough to produce it. A sufficient cause provides, or brings together, everything required for the effect to occur.

For example, consider a simple electrical lamp. Connecting an operational lamp to a suitable electrical supply through a complete circuit is sufficient for it to illuminate, provided that no other condition prevents its operation. The electrical transfer alone is necessary but not sufficient: the recipient process also depends upon the required configuration and other necessary causal inputs: the lamp, connections and appropriate circuit conditions are also required. In complex systems, what we describe as a single sufficient cause is therefore often a causal organisation containing several contributing causes.

This leads directly to the consideration of complex causality in the next section.

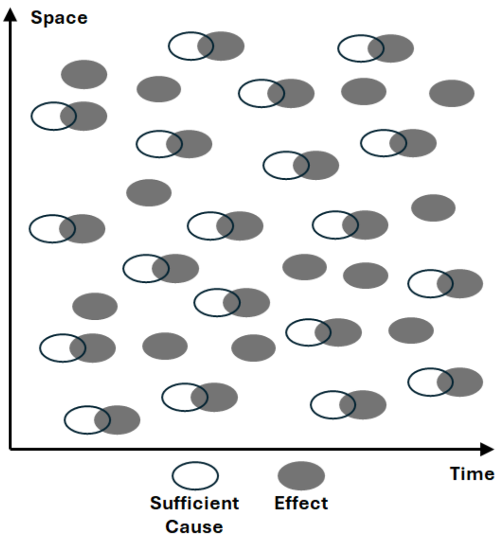

Figure 4. Sufficient Cause.

Conventionally, a sufficient cause is one in the presence of which the specified effect must occur (Mackie, 1974).

4.4 Complex Causality

Necessity is associated most fundamentally with the recipient process and what is transferred to it. However, the source of a necessary transfer is a separate question. Different source processes may be capable of producing equivalent transfers. For example, a required transfer of thermal energy might be supplied by electrical resistance, combustion or solar radiation. The recipient process need not depend upon which source provides the energy if each can provide the transfer required.

A particular source can therefore be described as necessary only where the required transfer cannot be supplied by other means. If source A is the only available source of a necessary transfer, then source A is necessary within the specified causal organisation and its removal prevents the recipient process from operating.

Where several alternative sources can provide the same necessary transfer, none of those sources is individually necessary for the resulting effect. Instead, they form alternative means of providing the necessary input. A particular source may therefore be unnecessary whilst the transfer that it provides is necessary.

This distinction also helps illuminate Rothman's sufficient-component cause model (Rothman, 1976). Different combinations of component causes may provide alternative sufficient pathways to the same effect. When these pathways are causally decompressed, some apparently alternative causes may prove to be different upstream sources capable of supplying the same necessary transfer. The apparent alternatives at the level of source causes need not therefore imply that the causal requirements of the recipient process themselves differ.

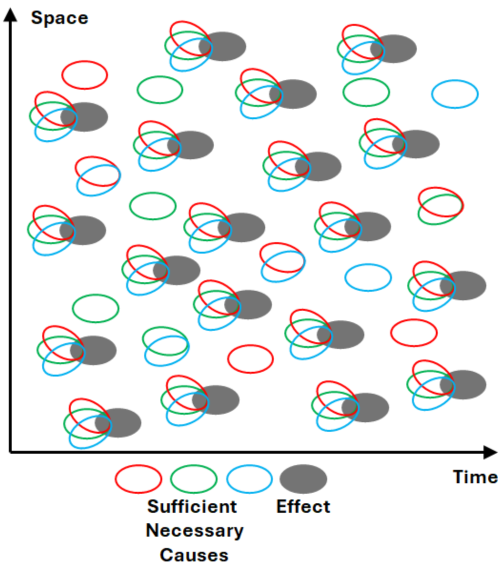

Figure 5. Necessary Causes Combining to Form a Sufficient Cause.
Where each necessary transfer has only one possible source, each source is necessary. Together, the three sources provide the complete set of necessary transfers and are sufficient for the specified effect.

4.5 Enabling and Inhibiting Causes

The preceding discussion of necessary and sufficient causes has principally concerned enabling causation: the causal inputs required to enable a process to function and produce an effect. A process may require several enabling inputs, each of which is necessary, while their complete combination may be sufficient to enable the process to operate.

However, causal influences may also be inhibiting. An inhibiting cause prevents, restricts or disrupts a process that would otherwise be capable of operating. Enabling and inhibiting causes therefore act in opposing ways: enablers make causal activity possible, while inhibitors prevent or limit that activity.

For example, a motor may have all of the electrical and mechanical inputs necessary for operation. A safety interlock may nevertheless prevent it from running. The enabling requirements have been satisfied, but an inhibiting causal pathway prevents the enabled process from operating.

Inhibiting causes, like enabling causes, may be considered in terms of necessity and sufficiency. A single inhibitor may be sufficient to prevent a process from operating, or several inhibiting causes may need to act together before inhibition is sufficient. Conversely, a particular inhibitor is necessary for an inhibitory effect where that effect cannot occur in its absence.

This reveals an important asymmetry between enablement and inhibition. Sufficient inhibition takes causal precedence over sufficient enablement. Even where the complete set of enabling inputs is present, a sufficient inhibiting cause can prevent or limit the process from functioning. Removing the inhibitor does not itself produce the resulting effect; it allows the enabling causal organisation to operate.

A process and its effect therefore depend upon the combined causal organisation of enabling and inhibiting influences. Explaining why something occurred requires understanding the enabling causal inputs that allowed the relevant processes to operate, while explaining why something did not occur may require identifying either missing enablers or effective inhibitors.

This distinction also provides a connection with the systems concepts of conditions and constraints developed later in the paper. Conditions can enable causal processes, while constraints may inhibit or restrict them or restrict the pathways through which their causal transfers occur.

4.6 Probabilistic Causation

Causal relationships are often described as probabilistic when the presence of a particular cause increases or decreases the likelihood of an observed effect of a particular type (Hitchcock, 2018). However, this does not necessarily mean that the causal influence itself is inherently probabilistic.

Consider a type of effect that requires several necessary causes acting together. If one necessary cause is known to be present but the presence of the others is uncertain, the known cause may appear to increase the probability of the effect. It can therefore be described epistemically as a probabilistic enabler, even though ontologically it remains a causal component within a larger causal organisation.

For example, adequate rainfall may be necessary for the successful growth of a particular crop. Nevertheless, observing adequate rainfall does not guarantee a successful harvest because suitable temperature, soil conditions, nutrients, absence of damaging disease and other necessary causes may also be required. If some of these are unknown, rainfall will appear statistically to increase the probability of a successful harvest.

The same principle applies to inhibitors. An identified inhibitor may reduce the observed probability of an effect because its effectiveness depends upon other causal circumstances whose presence is uncertain. What appears to be probabilistic inhibition may therefore result from incomplete knowledge of the causal organisation within which the inhibitor operates.

Probabilistic descriptions consequently often express epistemic uncertainty about a wider causal network rather than identifying a fundamentally different form of causation. Statistical relationships can tell us that the presence of a factor changes the observed likelihood of an outcome, while further causal investigation seeks to identify the additional enablers, inhibitors and pathways responsible for that variation.

This distinction does not require the assumption that all causation in reality is deterministic. It makes the more limited claim that an observed probabilistic relationship should not, by itself, be taken as evidence that the underlying causal relationship is intrinsically probabilistic.

5. Causal Chains and Networks

5.1 From Relationships to Chains

Individual causal relationships rarely occur in isolation. As discussed earlier, an effect produced by one causal interaction may subsequently act as a cause in another. In this way, causal relationships concatenate to form causal chains.

For example, heavy rainfall may increase river flow, increased flow may raise the river level, and the rising level may eventually cause water to overtop the riverbank. Each stage results from causal processes and transfers and provides causal inputs to the next.

A causal chain can therefore be represented simply as:

A → B → C → D

However, this representation is a useful compression of a more detailed causal reality. Each intermediate element occupies two causal roles: it is an effect of what precedes it and a cause of what follows.

Causal chains also allow effects to occur far from their original causes in both space and time. The original causal influence does not need to act directly upon the final effect. Instead, causal continuity is maintained through the sequence of intermediate processes and transfers. This is the basis of indirect causation.

Simple chains are nevertheless only one possible form of causal organisation. A cause may contribute to several subsequent effects, while an effect may depend upon several preceding causes (Pearl, 2009). Causal chains therefore branch and converge, producing the more extensive causal networks considered next.

5.2 Causal Networks

Causal chains become causal networks when causal pathways branch or converge. A single cause may contribute to several subsequent effects, while a particular effect may depend upon several preceding causes. In complex systems, both commonly occur together.

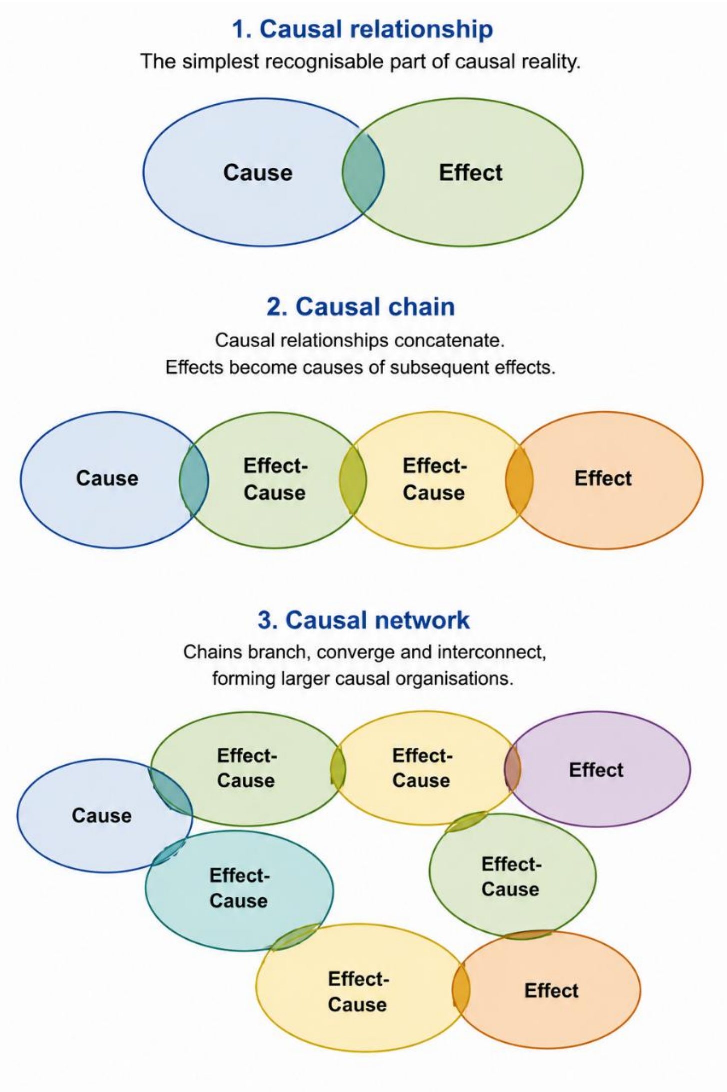

Figure 6. From Causal Relationship to Causal Network.
Individual causal relationships may concatenate into chains, while branching and convergence produce larger causal networks. The overlapping ellipses represent the continuity of causal relationships without yet decomposing them into their underlying processes and transfers.

Branching occurs when a causal influence produces consequences through more than one pathway. Heavy rainfall, for example, may increase river flow, saturate soils and raise groundwater levels. Each of these effects may then contribute to further effects of its own.

Convergence occurs when several causes contribute to the same effect. Flooding at a particular location might result from a combination of high river flow, saturated ground, restricted drainage and high downstream water levels. As discussed in Section 4, some of these causes may be individually necessary while becoming sufficient only when combined.

Real causal networks can therefore become extremely complex. They may contain many branching and converging pathways, with effects becoming causes of further effects. Some pathways may reinforce one another, while others may be inhibiting or constraining.

A causal network is consequently not simply a collection of independent cause-and-effect relationships. It is an interconnected causal organisation in which changes occurring in one part may propagate through multiple pathways and influence other parts of the network.

Representing causality as a network is particularly important in systems theory because it allows apparently separate effects to be traced back to shared causes and apparently simple effects to be understood as outcomes of multiple interacting causal pathways.

5.3 Direct and Indirect Causation

A distinction can be made between direct and indirect causation according to the causal pathway connecting a cause and effect (Dowe, 2000).

Direct causation occurs when a cause influences an effect through an identifiable causal interaction without an intervening causal system or event. For example, one moving billiard ball striking another transfers energy directly to it and changes its motion.

Indirect causation occurs when the causal influence passes through one or more intermediate processes or systems before producing the effect being considered. Each intermediate effect becomes a cause within the continuing causal chain.

For example, heavy rainfall in a river catchment may eventually cause flooding many kilometres downstream. The rainfall does not interact directly with the flooded buildings. Instead, it contributes to a sequence of intermediate processes involving runoff, river flow, rising water levels and overtopping. The relationship between the rainfall and the eventual damage is therefore indirect, while the complete causal pathway consists of a chain of more immediate causal interactions.

Indirect causal pathways also provide opportunities for other causal influences to intervene. At each intermediate stage, the continuing pathway may be enabled, inhibited, redirected or interrupted. Such intervention may be agentic, as when a person deliberately changes a process, or may result from interaction with another causal network.

For example, the causal pathway from heavy rainfall to downstream flooding may be modified by a reservoir retaining water, flood defences containing increased river levels, or deliberate operation of control structures. The greater the number of intermediate causal interactions, the greater the number of opportunities for such intervention.

Indirect causation may therefore be associated with greater uncertainty about whether an initial cause will ultimately produce a particular downstream effect. This does not necessarily make the underlying causation intrinsically probabilistic. Rather, additional intermediate interactions increase the number of causal circumstances about which knowledge may be incomplete.

The distinction between direct and indirect causation consequently depends partly upon the scale and boundary of analysis. A relationship described at a broad level as direct may reveal several intermediate causal interactions when examined more closely.

Indirect causation does not imply weaker or less genuine causation. Provided that a continuous pathway of processes and transfers connects the initial cause with the eventual effect, causal influence can propagate through many intermediate stages. Identifying those stages provides a more detailed explanation of how the effect was produced.

5.4 Causal Boundaries and Levels

Causal networks possess physical organisation arising from their constituent processes, transfers and persistent configurations. However, their division into particular causes and effects depends upon how that causal organisation is represented. An observer may select and bound portions of a larger causal network for investigation or explanation and attribute to them the relational roles of cause and effect. These epistemic boundaries do not create the underlying causal organisation, but determine which parts of it are represented as causes and effects for a particular purpose.

Consider the sequence:

Rainfall → runoff → river flow → rising water level → flooding → property damage

If the purpose is to explain increased river flow, rainfall and runoff may be treated as causes and increased flow as the effect. If the purpose is to explain property damage, flooding may instead be treated as the immediate cause, while rainfall becomes a more distant cause. The underlying causal network has not changed; only the part selected for attention has changed.

Causality can also be represented at different levels of detail. The statement *“heavy rainfall caused flooding”* compresses many intermediate processes and transfers into a single causal relationship. At a finer level of analysis, the same relationship may be decompressed into rainfall, infiltration, runoff, drainage, river flows, water levels and overtopping. Each of these may itself be decomposed further if greater causal detail is required.

The appropriate causal boundary and level of representation therefore depend upon the purpose of the investigation. A broad representation may be sufficient for some explanations, while diagnosis or intervention may require the causal pathway to be decomposed until the relevant processes and transfers can be identified.

Cause and effect should consequently be understood partly as epistemic selections from a larger causal reality. The observer determines which portion and level of the causal network is represented, while the underlying causal processes and transfers exist independently of that representation.

This distinction between causal reality and its representation prepares the ground for the two complementary PTP and TPT perspectives considered in the following section.

6. Processes, Transfers and Causal Duality

6.1 The Alternating Structure of Causality

The preceding sections have described causality in terms of processes connected by physical transfers. When these causal interactions are extended, they form an alternating sequence:

… P → T → P → T → P → T → P …

where P represents a process and T a transfer between processes. This alternating process–transfer structure provides the underlying ontological basis of the causal framework developed here.

A process receives matter, energy or information through one or more transfers. Its operation may then produce further transfers that become inputs to subsequent processes. For example, rainfall transfers water to the ground; hydrological processes generate runoff; runoff transfers water to a river; river processes convey the water downstream; and further transfers may eventually produce flooding. At greater levels of detail, each of these processes can itself be represented as further processes and transfers.

The P–T–P–T sequence should not be interpreted as implying that causal reality consists only of simple linear chains. Processes may receive several transfers and produce several others, causing causal pathways to converge and branch. The same alternating relationship between processes and transfers nevertheless remains present within each pathway through the resulting causal network.

Nor are P and T intended to represent two different kinds of causality. They represent the two alternating components through which causal continuity occurs: processes transform or act upon inputs, while transfers connect processes across space and time.

Earlier sections established that cause and effect are relational roles rather than fundamentally different kinds of causal phenomenon. What is represented as an effect in relation to preceding causality may subsequently be represented as a cause in relation to what follows. We should therefore not expect causes and effects to possess fundamentally different causal structures. If the same causal phenomenon can occupy either role, its underlying causal organisation does not change merely because its relational position within the causal representation changes.

This provides an important requirement for any representation derived from the P–T–P–T sequence: the same structural form should be capable of representing both causes and effects. Cause and effect should differ according to their relationship to preceding and subsequent causality, rather than according to their internal causal structure.

The alternating P–T–P–T structure satisfies this requirement in two complementary ways. By placing conceptual boundaries around different portions of the same sequence, causal phenomena can be represented as Process–Transfer–Process (PTP) or Transfer–Process–Transfer (TPT). In the PTP perspective, both causes and effects can be represented as events involving interaction between processes. In the TPT perspective, both can be represented in terms of the functioning, altered functioning or non-functioning of processes through their inputs and outputs.

PTP and TPT therefore do not represent different kinds of causality. They are complementary epistemic perspectives upon the same underlying P–T–P–T causal organisation. Their respective interpretations are developed in the following sections.

6.2 PTP — The Event Perspective

One way of representing the alternating causal sequence is to place a conceptual boundary around two processes and the transfer connecting them:

P → T → P

This is referred to as the Process–Transfer–Process (PTP) perspective. It foregrounds the interaction through which one process influences another and can therefore be understood as an event perspective.

For example, when one moving billiard ball strikes another, processes associated with the movement of the first ball result in a transfer of energy and momentum to the second, changing its motion. The PTP representation directs attention towards the transfer connecting the processes and therefore towards the mechanism through which the interaction occurs.

PTP is particularly useful when investigating whether an apparent causal relationship is physically connected. Rather than observing only that one change precedes another, the investigator asks what was transferred and through what pathway. Identifying that transfer provides evidence for the causal interaction.

The boundaries of a PTP event are nevertheless epistemic. The two processes and their connecting transfer remain parts of the larger alternating causal network:

… T → [P → T → P] → T → P …

PTP should therefore not be understood as a distinct form of causality existing within reality. It is a way of representing part of the underlying P–T–P–T causal organisation that foregrounds interaction, transfer and change.

6.3 TPT — The System Perspective

The same alternating causal sequence can be represented by placing a conceptual boundary around a process together with its incoming and outgoing transfers:

T → P → T

This is referred to as the Transfer–Process–Transfer (TPT) perspective. It foregrounds the process that transforms an input into an output and can therefore be understood as a system perspective (von Bertalanffy, 1968).

For example, consider a water pump. Water and energy enter the pump through causal transfers, processes within the pump increase the pressure of the water, and water at higher pressure leaves as a further transfer. Representing this as TPT directs attention towards the organised process through which inputs are transformed into outputs.

TPT is particularly useful when causal relationships between processes are recognised but the transfers connecting them are not fully known. We may observe that the functioning, altered functioning or non-functioning of one process is regularly associated with changes in another and infer that some causal transfer connects them without yet knowing its physical nature or pathway. The TPT perspective allows such causal relationships to be represented while leaving the intervening transfers provisionally unspecified, thereby providing a basis for hypotheses that can subsequently be investigated and causally decompressed.

As with PTP, however, the boundary is epistemic. The process and its input and output remain parts of the same larger causal network:

… P → [T → P → T] → P → T …

TPT is therefore not a second form of causality. It is an alternative representation of the same underlying P–T–P–T causal organisation, foregrounding the system and the transformation of inputs into outputs, rather than the interaction between adjacent processes.

6.4 Causal Duality

PTP and TPT provide two complementary ways of representing the same underlying causal reality. They do not describe two different forms of causation. The distinction between them arises from where an observer places conceptual boundaries within the continuing sequence of processes and transfers:

… P → T → P → T → P → T → P …

If the boundary encloses two processes and the transfer connecting them, the resulting representation is PTP. If it encloses a process together with its incoming and outgoing transfers, the representation is TPT. The underlying causal organisation remains unchanged. The distinction is therefore epistemic rather than ontological.

In either perspective, causal chains are formed from successive causes and effects having the same structure. This follows from the relational nature of cause and effect established earlier: what is represented as an effect in relation to preceding causality may subsequently be represented as a cause in relation to what follows. Continuity between successive cause-and-effect units is provided by a component common to adjacent units.

In the PTP perspective, causes and effects are PTP units and adjacent units share a process:

$[\mathbf{P_1} \rightarrow \mathbf{T_1} \rightarrow \mathbf{P_2}] \rightarrow [\mathbf{P_2} \rightarrow \mathbf{T_2} \rightarrow \mathbf{P_3}]$

The process $\mathbf{P_2}$ is common to both units. It is the receiving process in the first interaction and the originating process in the next.

PTP also provides the most direct representation of necessary and sufficient causation. Processes may require several causal transfers in order to function, so multiple causal pathways may converge upon a recipient process. Each required transfer may be necessary, while the complete set of necessary transfers may together be sufficient for the process to function in the absence of sufficient inhibition. Necessity and sufficiency therefore attach most directly to the convergence of causal transfers upon processes represented through PTP. As discussed earlier, the sources of those transfers need not themselves be necessary where alternative sources can provide equivalent transfers.

In the TPT perspective, causes and effects are TPT units and adjacent units share a transfer:

$[\mathbf{T_1} \rightarrow \mathbf{P_1} \rightarrow \mathbf{T_2}] \rightarrow [\mathbf{T_2} \rightarrow \mathbf{P_2} \rightarrow \mathbf{T_3}]$

Here the transfer $\mathbf{T_2}$ is common to both units. It is the output from one process and the input to the next. TPT consequently foregrounds the functioning of a process: given particular inputs, the process functions, functions differently or fails to function, with corresponding consequences for its outputs.

The transfer connecting successive processes may be difficult to identify or observe. We may therefore recognise that changes in one process are regularly followed by changes in another without knowing the physical transfer connecting them. The resulting representation compresses the causal relationship into something resembling **Process A → Process B**, while the intervening transfer remains hidden. This may contribute to the impression that causality itself is an abstract or non-physical relationship between events. Within the framework developed here, however, the absence of an identified transfer represents a limitation of causal knowledge rather than evidence for the absence of a physical causal pathway. Identifying the intervening transfer therefore amounts to decompression of the causal relationship and provides a more complete causal explanation.

The two perspectives consequently divide the same alternating causal sequence differently. PTP foregrounds interactions between processes and the convergence of causal inputs upon them, while TPT foregrounds the functioning of processes as transformations of inputs into outputs. Neither changes the underlying causal reality; each provides a different epistemic decomposition of it.

This causal duality provides a bridge between systems and causality. Systems can be examined through the TPT perspective as processes transforming inputs into outputs, while their interactions can be examined through the PTP perspective as processes connected by transfers. Moving between the two therefore allows the same causal network to be understood alternately in terms of process functioning and causal interaction.

TPT representation is well suited to the recognition of recurring patterns in process behaviour. Inputs, process functioning and outputs can recur in recognisable configurations without requiring the observer to possess a complete causal account of the wider network in which they occur. Such patterns may provide useful predictions and generate hypotheses about underlying causality.

PTP representation, by contrast, facilitates the logical analysis of causal networks. It directs attention towards the transfers connecting processes, the convergence and branching of causal pathways, the necessity or sufficiency of converging causal contributions, alternative sources of required transfers, and the presence of enabling or inhibiting causal pathways. Moving from recognised patterns of process behaviour towards identification of these relationships therefore amounts to causal decompression and can provide a more complete causal explanation.

The two perspectives may consequently support different forms of causal cognition. TPT facilitates recognition of recurring patterns between the behaviour of processes, whereas PTP facilitates explicit logical analysis of the causal networks through which those behaviours arise. This possible relationship between causal representation and human cognition is considered further in subsequent work.

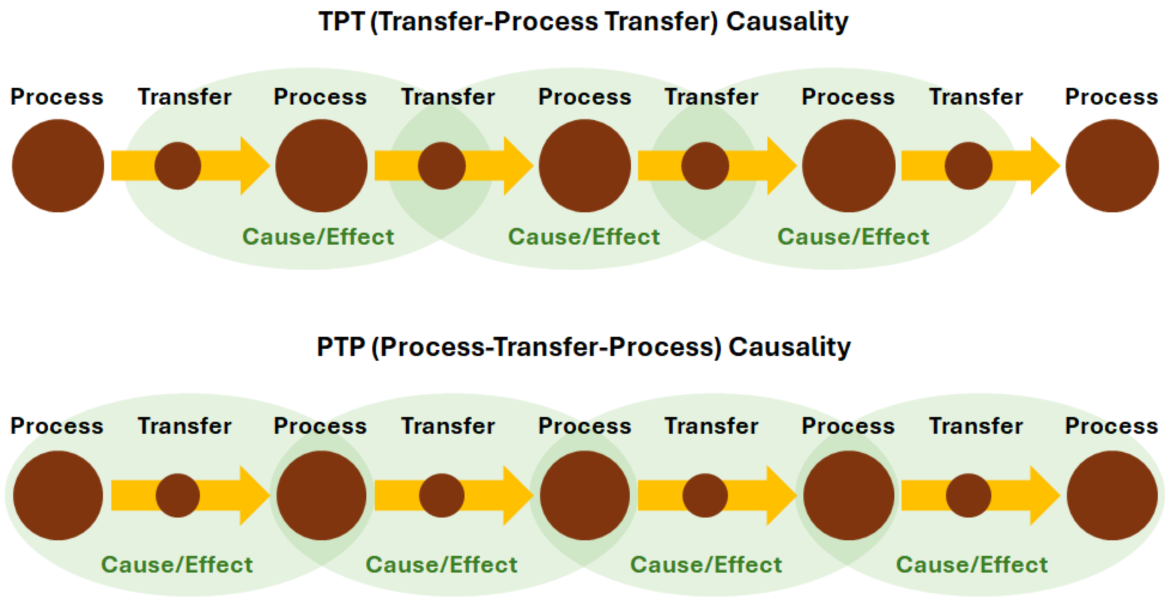

Figure 7. Causal Reality and its PTP and TPT Representations

The underlying P–T–P–T causal sequence remains unchanged; only the observer's conceptual boundary moves. PTP provides an event-based perspective in which interactions between processes are represented as causes and effects, while TPT provides a systems-based perspective in which process functioning is represented through inputs and outputs. The two are complementary representations of the same causal reality: PTP foregrounds causal interaction and network structure, while TPT foregrounds system functioning. Together they provide a conceptual bridge between causal theory and systems theory.

7. Causal Structures

7.1 Recurring Causal Organisation

Causal networks may contain recurring patterns in the ways that causes and effects are organised. Some causal pathways proceed sequentially, while others branch, converge, recur or return causal influence to earlier parts of the network. When such patterns recur, they can be recognised and represented as causal structures.

Causal structures are therefore epistemic representations of recurring organisation within underlying causal networks. Recognising them allows complex causal relationships to be described more simply and compared across otherwise very different systems. The principal causal structures considered within this framework are introduced in the following sections.

7.2 Linear Causality

Linear causality is the simplest causal structure, in which a cause produces an effect through a single causal pathway. For example, heating a metal rod causes it to expand.

Linear representations are useful for simple causal relationships, although they may compress wider causal networks that are not themselves linear.

7.3 Causal Trees and Roots

A causal tree occurs when one cause gives rise to several effects through branching causal pathways. For example, a volcanic eruption may produce lava flows, ash fall and atmospheric effects.

A causal root has the opposite structure: several causal pathways converge upon a single effect. For example, river flooding may result from the combined effects of heavy rainfall, saturated ground, high groundwater and restricted drainage.

7.4 Domino Causality

Domino causality is a sequential causal structure in which each effect becomes the cause of the next effect in the sequence. A row of falling dominoes provides the literal example: each falling domino transfers energy to the next, causing it to fall in turn.

7.5 Cyclic Causality

Cyclic causality occurs when a sequence of causal relationships repeatedly recreates the same types of causes and effects. The individual entities or events need not be the same on each cycle.

For example, a chicken produces an egg from which another chicken develops, which may subsequently produce further eggs. The causal pattern recurs even though each chicken and egg is a different entity.

7.6 Mutual Causality

Mutual causality occurs when two entities or processes causally influence one another.

For example, two people engaged in conversation continually influence each other's subsequent responses. Each is therefore simultaneously part of the causal environment of the other.

7.7 Feedback

Feedback occurs when an output from a process or system returns, directly or indirectly, as an input that influences the subsequent operation of that process or system (Wiener, 1948).

For example, a thermostat responds to room temperature and switches a heating system on or off. The resulting change in temperature is subsequently detected by the thermostat and influences its further operation. Feedback therefore forms a closed causal pathway through which the consequences of a process influence its continuation or modification.

7.8 Cascading Causality

Cascading causality is a causal chain in which each successive stage produces both the next stage in the chain and an additional contribution to a common effect. As the chain continues, the contributions to the common effect accumulate, potentially increasing its magnitude.

For example, a person experiencing difficulty repaying debt may borrow additional money, creating a further repayment obligation that contributes to further borrowing. Each stage continues the causal chain while also increasing the person's total indebtedness. This produces the familiar effect of progressively “digging oneself into a hole”.

Cascading causality therefore differs from simple domino causality because successive stages do more than propagate the causal sequence: they progressively contribute to a common accumulating effect.

7.9 Propagating Causality

Propagating causality occurs when causal effects spread through a connected network, producing further effects that initiate additional causal pathways elsewhere in the network. The number of affected components or processes may therefore increase as the causal influence propagates.

For example, failure of one component in an electricity network may increase the load upon neighbouring components. Some of these may then fail, transferring additional loads to other components and causing the disturbance to spread through the network (Motter & Lai, 2002).

Propagating causality differs from cascading causality because its defining feature is the spread of causal effects through a network, rather than successive stages contributing to a common accumulating effect. It also differs from feedback because the causal influence need not return to its source.

7.10 Relational and Threshold Causality

Relational causality occurs when an effect depends upon a relationship or difference between the states of two or more entities rather than upon the state of either considered alone. For example, heat flows between two objects because of the difference in their temperatures.

Threshold causality occurs when an effect is produced only when a relevant variable reaches or crosses a particular value. For example, water begins to overflow from a container when its level rises above the height of the containing edge. Below that threshold, increasing the water level does not produce overflow; once the threshold is crossed, a new causal pathway becomes active.

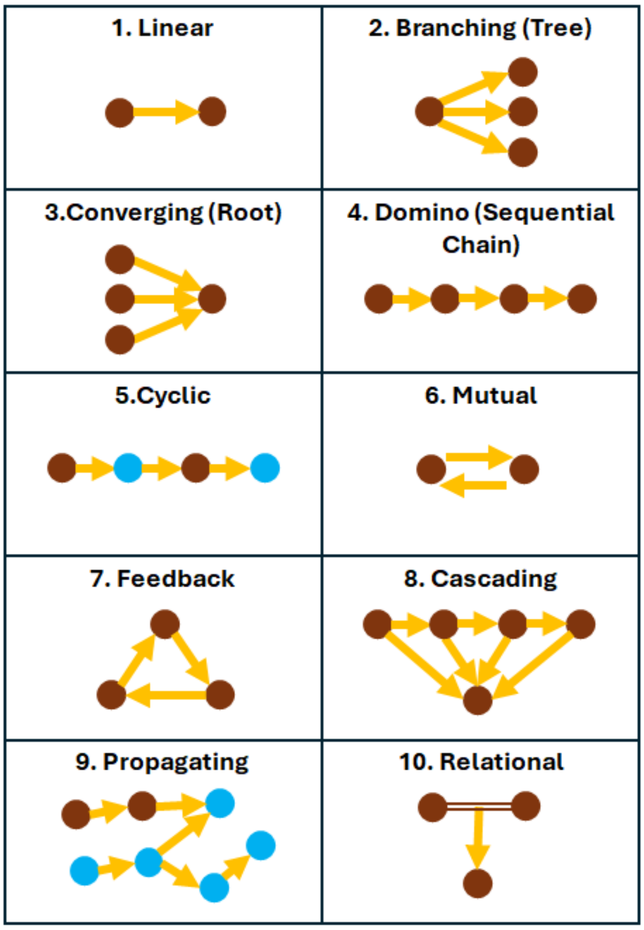

Figure 8. Common Causal Structures

8. Persistent Causal Organisation

8.1 Transient and Persistent Causal Networks

Causal networks differ greatly in their duration. Some form briefly around a particular event and then cease to exist, while others persist through time and continually influence subsequent events. It is therefore useful to distinguish between transient and persistent causal networks.

A transient causal network is associated with a particular causal episode. For example, a vehicle collision involves a network of interacting processes and transfers that develops over a short period and then largely comes to an end. Its consequences may persist, but the causal organisation associated with the collision itself is temporary.

A persistent causal network, by contrast, continues to operate over an extended period. A river system, for example, continually receives, stores and transfers water through a network of hydrological processes. The particular water and the states of individual components change, but the causal organisation persists.

Persistence therefore does not require the components of a causal network to remain unchanged. Matter, energy and information may continually flow through it, individual processes may change, and components may be replaced while the broader causal organisation continues.

Persistent causal networks are particularly important because they form part of the causal context within which transient events occur. They continually influence which causal pathways are available, which processes can operate, and which effects are possible. In this sense, many of the circumstances, conditions and constraints within which events occur can themselves be understood as persistent causal networks rather than as passive background features.

8.2 Circumstances, Conditions and Constraints

Events and processes occur within surrounding circumstances that are commonly treated as background features of the situation. From a causal perspective, however, many such circumstances can be understood as persistent causal networks that continue to operate while more transient causal events occur. They may therefore interact with those events and influence which causal possibilities can be realised.

For example, a road journey occurs within circumstances that include the road network, other traffic, weather and the prevailing physical environment. These are not simply descriptive features: each arises from continuing causal processes that may influence the journey without necessarily constituting its immediate cause.

Circumstances may have different causal relationships with particular processes. A condition is a circumstance that enables a process to occur, while a constraint is a circumstance that restricts or inhibits a process or channels the pathways through which its causal transfers can occur. Sufficient oxygen, for example, provides a condition necessary for ordinary combustion, while the banks of a river constrain and channel the movement of water.

Conditions and constraints are therefore relational rather than fundamentally different kinds of causal organisation. The same persistent causal network may enable one process while constraining another. A canal, for example, constrains the lateral movement of water while thereby enabling its controlled conveyance along a particular route.

What is treated as a circumstance, condition or constraint also depends partly upon the purpose and boundary of the investigation. A causal network treated as part of the surrounding circumstances in one analysis may itself become the principal object of investigation in another.

Circumstances should therefore not be regarded merely as passive backgrounds against which causation occurs. They may be persistent causal networks that interact with other causal networks, enabling, inhibiting, restricting or channelling their operation. Conditions and constraints identify particular causal relationships between those circumstances and the processes under consideration.

8.3 Configuration, Persistent Configuration and Causality

The configuration of entities influences causality because their relative arrangement affects which interactions are possible and the pathways through which causal transfers can occur. Configuration therefore helps establish the potential causal pathways available at a particular time.

For example, the configuration of components in an electrical circuit determines which components are connected and therefore the pathways through which electrical transfers may occur. Similarly, the configuration of pipes, valves and reservoirs within a water supply network determines the possible pathways through which water can flow.

Potential causal pathways should be distinguished from their realisation. A pipe may connect two tanks without water flowing between them. When an appropriate pressure difference exists and a valve is open, the potential pathway is realised through an actual transfer.

Where a configuration forms a persistent organised arrangement that supports recurring causal relationships, it may be recognised as a persistent configuration. Persistent configuration can therefore embody relatively persistent causal possibilities while particular processes and transfers realise those possibilities at particular times.

Configuration and structure also constrain causality by making some interactions possible while excluding others. A change in configuration may consequently alter the causal possibilities of a system even before any new causal pathway is realised.

Configuration establishes possibilities for causal interaction; structure represents persistent organised configurations of those possibilities; processes and transfers realise particular causal pathways.

Figure 9. illustrates the difference between potential and realised causal pathways in a persistent causal organisation. However, it should be noted that persistent causal organisation need not consist solely of inactive pathways awaiting transient activation. Some causal interactions may themselves persist, with transfers occurring continuously or repeatedly between processes. A persistent causal organisation may therefore contain both persistent causal activity and pathways that are activated only transiently. The particular behaviour of a system at any time arises from the combination of these continuing interactions and the transient causal activity occurring through its wider persistent configuration.

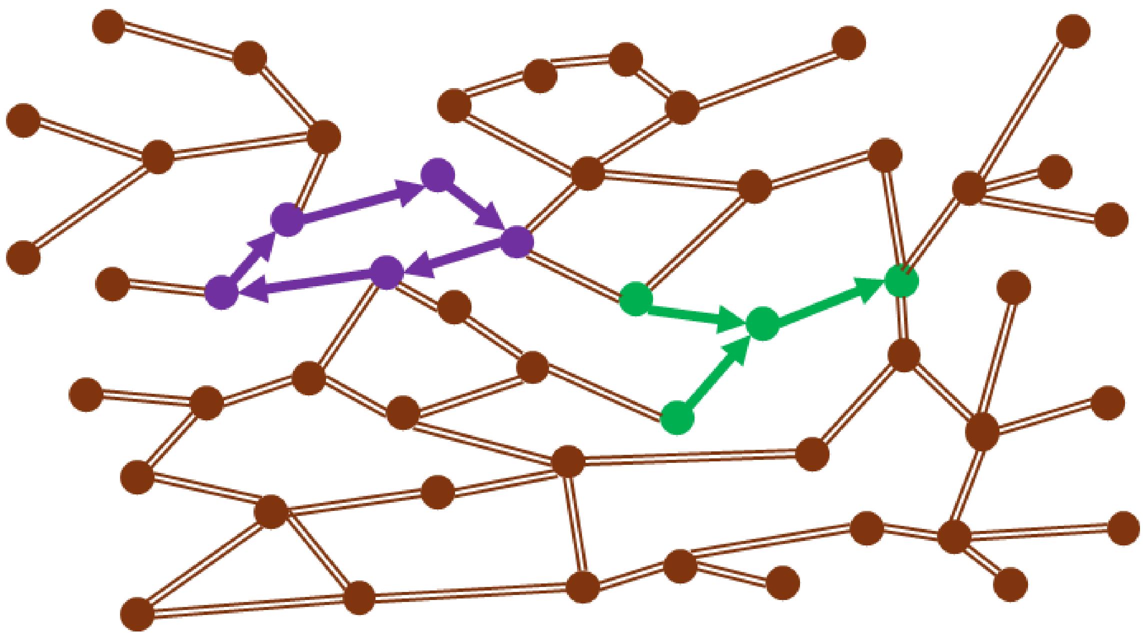

Figure 9. Transient Causal Activity within Persistent Causal Organisation.
The persistent organisation of a system establishes causal pathways through which transfers may occur. At any particular time, only some of these possibilities may be realised as transient causal activity. Different patterns of transient activity may occur simultaneously within the same persistent organisation and constitute aspects of the system's behaviour.

9. Causal Organisation and Systems

9.1 Systems as Persistent Causal Organisations

A system can be understood as an organised concentration of causal relationships. Its components do not merely coexist within a boundary; they interact through processes and transfers that together produce the properties, behaviours or outputs associated with the system as a whole (von Bertalanffy, 1968).

For example, a motor consists of components arranged so that electrical, magnetic and mechanical interactions occur in an organised manner. These causal relationships transform electrical inputs into mechanical outputs. A collection containing the same components but lacking the required organisation would not function as a motor.

The causal relationships within a system are also typically more concentrated than those connecting it with its wider environment. Components interact repeatedly with particular other components, creating persistent causal pathways through which matter, energy and information are transferred and transformed.

This does not mean that systems are causally isolated. They remain connected to other systems and to their environments through inputs and outputs. Rather, the concentration and organisation of causal relationships provide part of the basis upon which a portion of a larger causal network can be recognised and represented as a system.

Systems can therefore be viewed not simply as collections of components, but as persistent causal organisations in which interacting components and processes are organised sufficiently to exhibit emergent properties, behaviours or outputs.

9.2 Processes as Realised Causality

As established in Section 8, persistent configurations can embody potential causality by maintaining pathways through which future causal interactions may occur. These possibilities become realised causality when processes occur and transfers take place.

For example, the configuration of a road network provides many possible pathways between locations. A particular journey realises one of those possibilities through the movement of a vehicle along a particular route. The other routes remain available as potential causal pathways but are not involved in that particular journey.

Processes realise causality by receiving inputs, acting upon or transforming them, and producing outputs that may become inputs to subsequent processes. Transfers connect these processes, allowing causal influence to propagate through the wider causal network.

The distinction between potential and realised causality is important because the same persistent configuration can support many different patterns of causal activity. Which possibilities are realised depends upon the conditions, constraints and causal inputs present at a particular time.

A system's causal organisation therefore includes both persistent configurations that make particular interactions possible and the processes and transfers through which particular causal possibilities are realised.

9.3 Differential Coupling and Near-Decomposability

Persistent configurations may establish different degrees of causal connectivity between components. Some components may be connected by pathways through which matter, energy or information can be transferred frequently, rapidly or strongly, while others may interact only weakly, indirectly or occasionally. These differences in the extent to which components are causally connected can be described as differential coupling.

Where causal coupling is substantially stronger among one group of components than between those components and their surroundings, that concentration of causal interaction may provide a basis for recognising the group as a system or subsystem.

For example, the persistent configuration of the components within a motor makes repeated electrical, magnetic and mechanical interactions between those components possible. These interactions are generally more concentrated than the motor's interactions with other components of the larger machine.

Differential coupling can therefore produce relatively distinct concentrations of causal activity within larger causal networks. Where internal interactions are substantially stronger, denser or faster than external interactions, the resulting causal organisation may behave with some degree of relative independence.

This provides a causal basis for near-decomposability. A complex causal organisation may be usefully decomposed into relatively distinct parts because interactions within those parts are more strongly coupled than interactions between them (Simon, 1962). The boundaries are not absolute: matter, energy and information may continue to cross them.

Differential coupling does not, by itself, establish that such a causal organisation is a system or subsystem. Rather, it helps identify relatively coherent causal organisations that may warrant representation as wholes. Whether such an organisation is recognised as a system also depends upon the emergence of properties, behaviours or outputs attributable to that organisation as a whole.

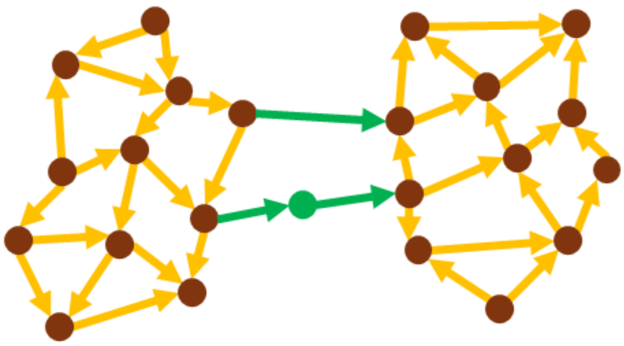

Figure 10. Differential Coupling and the Recognition of System Boundaries
Causal interactions may be more concentrated among some groups of processes than between those groups and their surroundings. Such differential coupling allows relatively coherent causal organisations to be recognised as systems or subsystems, even though causal interactions continue across their boundaries.

9.4 Emergence and System Recognition

Differential coupling can identify a relatively coherent causal organisation, but this alone does not necessarily make that organisation a system. A further consideration is whether the organised interactions produce properties, behaviours or outputs attributable to the organisation as a whole rather than to its components considered independently.

These are emergent properties, behaviours or outputs. For example, the components of an electric motor, when appropriately configured and causally interacting, produce rotational motion. No individual component possesses this behaviour independently; it arises from their organised causal interaction.

Emergence therefore provides an additional basis for recognising a relatively coherent causal organisation as a system. Persistent configuration establishes continuing possibilities for interaction, processes and transfers realise those possibilities, and differential coupling produces a relatively coherent causal organisation. Where that organisation exhibits emergent properties, behaviours or outputs when represented as a whole, it may be recognised as a system.

Once established, a system's causal organisation may itself contribute to maintaining the persistent configuration upon which its continued operation depends. In a living organism, for example, continuing biological processes repair components, regulate internal conditions and replace materials, thereby helping to maintain the configuration required for those processes to continue. Persistent configuration may therefore be both a prerequisite for, and an outcome of, continuing system operation.

Emergence may also create causal capabilities that the individual components do not possess independently. An organised system may consequently produce transfers and effects upon other systems that none of its components could produce alone. These system-level causal capabilities remain physically realised through the processes and transfers of the components, but arise only from their particular organisation.

Systems may themselves become components of larger systems whose organisation produces further emergent causal capabilities. Emergence can therefore generate an ontological hierarchy of causal organisation, in which causal capabilities arise at successive levels that are unavailable to the components at the level below when considered independently.

The same principle allows a relatively coherent causal organisation within a larger system to be represented as a subsystem where it exhibits its own emergent properties, behaviours or outputs while continuing to participate causally in the larger organisation.

System recognition can therefore be understood as arising from the conjunction of persistent configuration, differential coupling and emergence, together with the causal capabilities generated by their organisation. A fuller account of the processes through which emergence occurs is developed separately in a subsequent paper.

9.5 Assembly as an Effect

Assembly can be understood as a distinctive causal effect in which existing entities are transferred and brought together into a new configuration. Unlike many causal interactions, where only some matter, energy or information passes from a source to a recipient, in assembly the source entity itself may constitute the transfer in its entirety.

For example, during manufacture, a component may be moved from one location and incorporated into a motor. Relative to its original configuration, the component changes from existing there as an independent entity to no longer existing there. However, it has not ceased to exist absolutely: the whole entity has been transferred and now persists as a component of the resulting assembly. The source, transfer and resulting effect are therefore connected through continuity of the transferred entity.

Assembly consequently differs from an effect involving only a change in the state of an existing recipient. The causal processes bring previously separate entities into a new configuration, while those entities retain sufficient continuity to remain identifiable as components of the resulting assembly.

The new configuration may also establish causal possibilities that did not exist while the components were separate. Assembly does not necessarily produce a system, but where the resulting configuration persists, supports differential coupling between its components, and gives rise to emergent properties, behaviours or outputs, the assembly may be recognised as a system.

An assembly produced as an effect may subsequently become a cause within further causal relationships. An assembled motor, for example, can produce mechanical transfers capable of causing effects elsewhere in a machine. Assembly therefore illustrates the earlier principle that an effect may subsequently become a cause.

Assembly is also important in the development of ontological hierarchy. Existing entities may be assembled into systems, and those systems may subsequently become components of still larger assemblies. Where new levels of organisation produce emergent causal capabilities, assembly can contribute to successive levels of causal organisation.

Assembly can therefore be understood as a causal effect in which whole entities are transferred into a new configuration, where they may persist as components and establish new causal possibilities.

10. Causal Leverage

10.1 Unequal Causal Contributions

Where several causes contribute to an effect, their physical contributions need not be equal. One cause may provide most of the matter or energy involved, while another provides only a small transfer that enables, inhibits or redirects the larger causal process.

A simple example is the ignition of fuel. The fuel contains the chemical energy subsequently released, while a spark supplies only a very small amount of energy. Nevertheless, the spark can make a necessary causal contribution by initiating the process through which the much larger store of chemical energy is released.

It would therefore be misleading to judge the causal importance of a contribution simply by comparing the quantities of matter or energy transferred. A relatively small causal input may make the difference between a larger causal process occurring and not occurring.

The same principle applies where a small input alters the pathway taken by a larger flow. Moving a valve requires relatively little energy, but the resulting change in configuration may determine whether a much larger flow of water occurs and where that water goes. The energy producing the water flow originates elsewhere; the smaller causal contribution changes the possibilities available to it.

Causal contributions can therefore differ both in their physical magnitude and in their influence upon the resulting causal organisation. This asymmetry provides the foundation for causal leverage, in which relatively small causal inputs can influence much larger flows of matter or energy.

10.2 Causal Leverage through Enablement, Inhibition and Control

Causal leverage occurs where a relatively small causal input of matter energy or information modifies a process or constraint that regulates a larger independently available causal flow. The magnitude of the resulting effect need not therefore correspond to the matter or energy contained in the initiating transfer, because the process's larger causal resources originate elsewhere.

A valve provides a simple example. A relatively small force may open, close or partially restrict a pathway through which a much larger flow of water occurs. The force applied to the valve does not provide the energy responsible for the water flow; it changes the constraint governing the pathway through which that independently available flow can occur. Similarly, railway points can redirect a much larger moving train by changing the configuration of the pathway available to it.

The causal role of constraints in enabling, restricting and redirecting system behaviour has also been emphasised in complex-systems accounts of causation and agency (Juarrero, 1999).

Causal leverage may therefore operate through enablement, inhibition or control. A causal input may make a pathway available, prevent its use, alter its capacity or select between alternative pathways. In each case, a relatively small causal input changes the causal organisation governing a potentially much larger effect.

Information can provide a particularly important form of causal leverage. A relatively small transfer carrying information may change the configuration or operation of a process controlling a much larger flow. For example, a control signal may specify how far a sluice gate should be opened. The receiving control system responds to differences in the structure of the signal by changing the position of the gate and thereby regulating the much larger flow of water.

The causal influence of information depends not primarily upon the quantity of energy carrying it, but upon the ability of a receiving process to respond differently to different informational structures. Variation in the structure of the controlling transfer can therefore produce variation in the resulting causal effect without supplying the matter or energy responsible for that effect.

Causal leverage consequently provides a general physical explanation for how relatively small causes can have disproportionately large consequences. The initiating cause modifies the conditions, constraints or pathways governing larger causal resources rather than itself supplying those resources. Informational control is an important special case of this more general principle.

10.4 Causal Leverage and Agency

Causal leverage provides a physical basis through which agency can produce substantial effects upon the external environment. Within a human agent, relatively small transfers associated with neural activity can initiate and control the much larger energy transfers involved in muscular activity. The muscles can then act upon the external environment, thus altering processes, configurations, conditions or constraints beyond the agent:

Neural activity → muscular activity → external action → wider causal consequences

For example, muscular action may operate a valve controlling a much larger flow of water, operate a switch controlling machinery, introduce a small quantity of a chemical that substantially alters a biological or chemical process, or communicate information that changes the subsequent behaviour of another agent.

In each case, the initiating action need not itself supply the matter or energy involved in all of the subsequent consequences. Rather, it acts at a point within a causal network from which effects can propagate, processes can be initiated or inhibited, constraints can be altered, or independently available causal resources can be redirected.

The causal significance of agency therefore depends not simply upon the magnitude of the matter or energy directly transferred by the agent, but upon where and how the agent can intervene within wider causal networks. Causal leverage allows relatively small agentic actions to produce disproportionately large or extensive consequences through the causal organisation with which they interact.

Agency can therefore be understood physically as the capacity of an organised system to generate causal transfers that selectively influence subsequent causal pathways. Through causal leverage, an agent with limited immediate energy can produce effects involving much greater flows of matter and energy available elsewhere in the environment.

This provides a foundation for later treatments of agency. What appears at the level of the whole as an agent acting upon the world can be decomposed into an organised sequence of physical processes and transfers through which relatively small causal inputs progressively leverage larger causal flows.

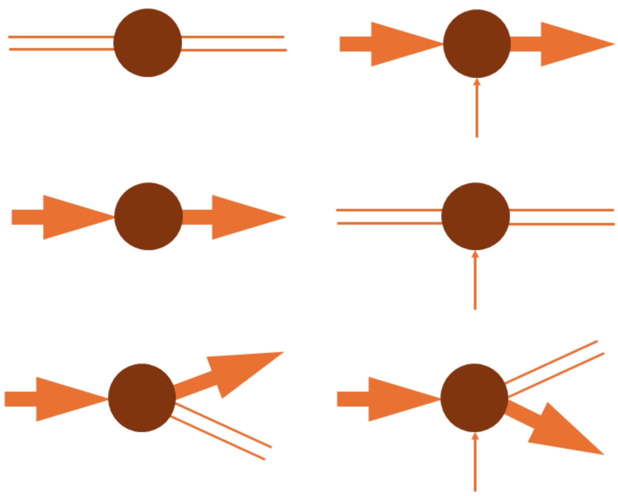

Figure 11. Causal Leverage through Enablement, Inhibition and Control.
A relatively small causal input may exert leverage by enabling an available causal pathway (top), inhibiting an active causal pathway (middle), or controlling causal activity by selecting between alternative available pathways (bottom). Causal leverage arises from the effect of the input upon causal organisation rather than simply from the magnitude of the initiating transfer.

11. From Causal Reality to Causal Knowledge

11.1 Recognition and Investigation of Causal Relationships

The preceding sections have principally considered causal reality: the processes, transfers and causal organisations through which effects occur. Our knowledge of causality, however, depends upon our ability to recognise possible causal relationships and investigate how they are physically instantiated.

An important starting point is the recognition of recurrence. When similar events repeatedly occur together or in sequence, or when variation in one is consistently associated with variation in another, we may recognise a pattern and suspect that a causal relationship exists. Such recurring associations may be described as correlations and can provide the basis for forming causal hypotheses.

Correlation does not, however, establish causation (Pearl, 2009). Two phenomena may recur together because both arise from a common cause, because the direction of causality has been misunderstood, because several causal pathways are involved, or because relevant conditions and constraints have not been identified. Increased ice-cream sales and cases of sunburn, for example, may correlate because warmer, sunnier weather contributes to both. Correlation therefore identifies a relationship requiring explanation rather than providing the explanation itself.

This distinction is particularly important because the complete causal pathway is often not directly observable. We may recognise a recurring relationship between processes while the transfers and intermediate processes connecting them remain unknown. Pattern recognition can therefore identify where causality may exist before its physical basis is understood.

Stronger causal explanation requires investigation of the causal mechanism connecting the phenomena (Machamer et al., 2000). Within the framework developed here, a causal mechanism consists of the processes, transfers, conditions and constraints through which causal influence propagates. For example, the recurring relationship between heavy rainfall within a catchment and rising river levels downstream can be explained by identifying the intervening transfers and processes through which rainfall reaches the land, generates runoff, enters streams and rivers, and is conveyed downstream.

Identification of a proposed mechanism also enables more rigorous testing. Measurements can determine whether predicted transfers occur and whether their timing and magnitude are consistent with the proposed causal pathway. Interventions, naturally occurring differences and other observations may provide further opportunities to test particular parts of the explanation. Complete knowledge of every process and transfer is not necessary, but confidence in a causal explanation can increase as more of the proposed causal organisation is identified and successfully tested.

Causal investigation can therefore be understood as a process of causal decompression. A recurring relationship initially represented simply as **A → B** can progressively be expanded to reveal the processes, transfers, conditions and constraints through which A contributes to B. Recognition of recurring patterns and correlations provides a means of generating causal hypotheses; identification and testing of the underlying causal organisation provides progressively stronger grounds for accepting or rejecting them.

11.2 Hypotheses, Explanations and Theories as Representations of Causal Reality

Hypotheses, explanations and theories can all be understood as epistemic representations of ontological causal reality. They represent selected processes, transfers, conditions, constraints and relationships within causal networks rather than constituting those networks themselves. Their boundaries and level of decompression depend upon the purpose of the representation and the knowledge available.

They differ principally in their epistemic status, purpose and scope. A causal hypothesis proposes a possible causal organisation that might account for an observed or anticipated phenomenon and can subsequently be tested against causal reality. An explanation represents causal organisation considered sufficiently well supported to account for how or why a phenomenon occurs. A theory extends such representation beyond particular instances by representing causal organisation believed to recur across a class of phenomena.

A hypothesis need not initially identify every process and transfer involved. Some parts of the proposed causal network may be known while others remain uncertain. Representing the hypothesis causally allows its components and relationships to be examined separately, so evidence may support some parts while contradicting others. Hypothesis testing can therefore be understood as comparing a proposed representation with observations and other evidence, actively seeking points at which it fails and progressively revising the representation where necessary (Popper, 1959).

Explanations similarly vary in their degree of causal decompression. The statement "heavy rainfall caused the flood" may provide an adequate explanation for one purpose, while another may require decomposition into rainfall, infiltration, runoff, river flow, drainage constraints, rising water levels and overtopping. An explanation is therefore always selective and may remain incomplete or partly incorrect as understanding develops.

Theories differ primarily in their wider scope. Plate tectonic theory, for example, does not merely explain one earthquake or volcanic eruption but represents recurring causal organisation associated with the movement and interaction of tectonic plates that can contribute to explanations of many geological phenomena. Where the represented causal organisation genuinely recurs, a theory may therefore provide both explanatory and predictive power.

These categories are not necessarily fixed. A representation initially proposed as a hypothesis may, following successful testing, contribute to an explanation; recognition that the same causal organisation recurs across many instances may subsequently contribute to theory. Conversely, observations inconsistent with an established explanation or theory may require its causal representation to be revised and its components tested as new hypotheses.

Hypotheses, explanations and theories therefore differ less in the kind of reality they represent than in the confidence placed in the representation, the purpose for which it is used and the range of phenomena to which it is intended to apply. Their development reflects a continuing correspondence between epistemology and ontology: causal organisations exist and recur within reality, while their recognition, representation and testing allow progressively more reliable causal knowledge to be constructed.

11.3 Explanatory Power and Recurrence

The explanatory power of a theory depends partly upon the extent to which the causal organisation it represents recurs in reality. A theory that represents a causal organisation found across many different instances can contribute to explaining a correspondingly wide range of phenomena.

For example, natural selection has substantial explanatory power because the causal relationships it represents recur across many populations, species and environments. The particular organisms and circumstances differ, but sufficiently similar causal organisation recurs for the same theoretical framework to contribute to their explanation.

This does not mean that a theory with wider application is necessarily better than one concerned with a narrower causal organisation. A narrowly applicable theory may provide a more detailed or precise explanation of a particular class of phenomena. Explanatory reach and explanatory adequacy are therefore different considerations.

Recurrence also helps explain why theories can be transferred between apparently different domains. If sufficiently similar causal organisation occurs in two different types of system, a theory developed to represent that organisation in one domain may also contribute to explaining phenomena in the other (von Bertalanffy, 1968). The surface characteristics of the systems may differ considerably while their underlying causal organisation is similar.

Explanatory power can therefore be understood partly as a relationship between an epistemic representation and the recurrence of the causal organisation it represents. The more widely that organisation genuinely recurs, the greater the potential explanatory reach of the theory. This also explains why identifying recurring causal organisation is so important to the development of general systems theory.

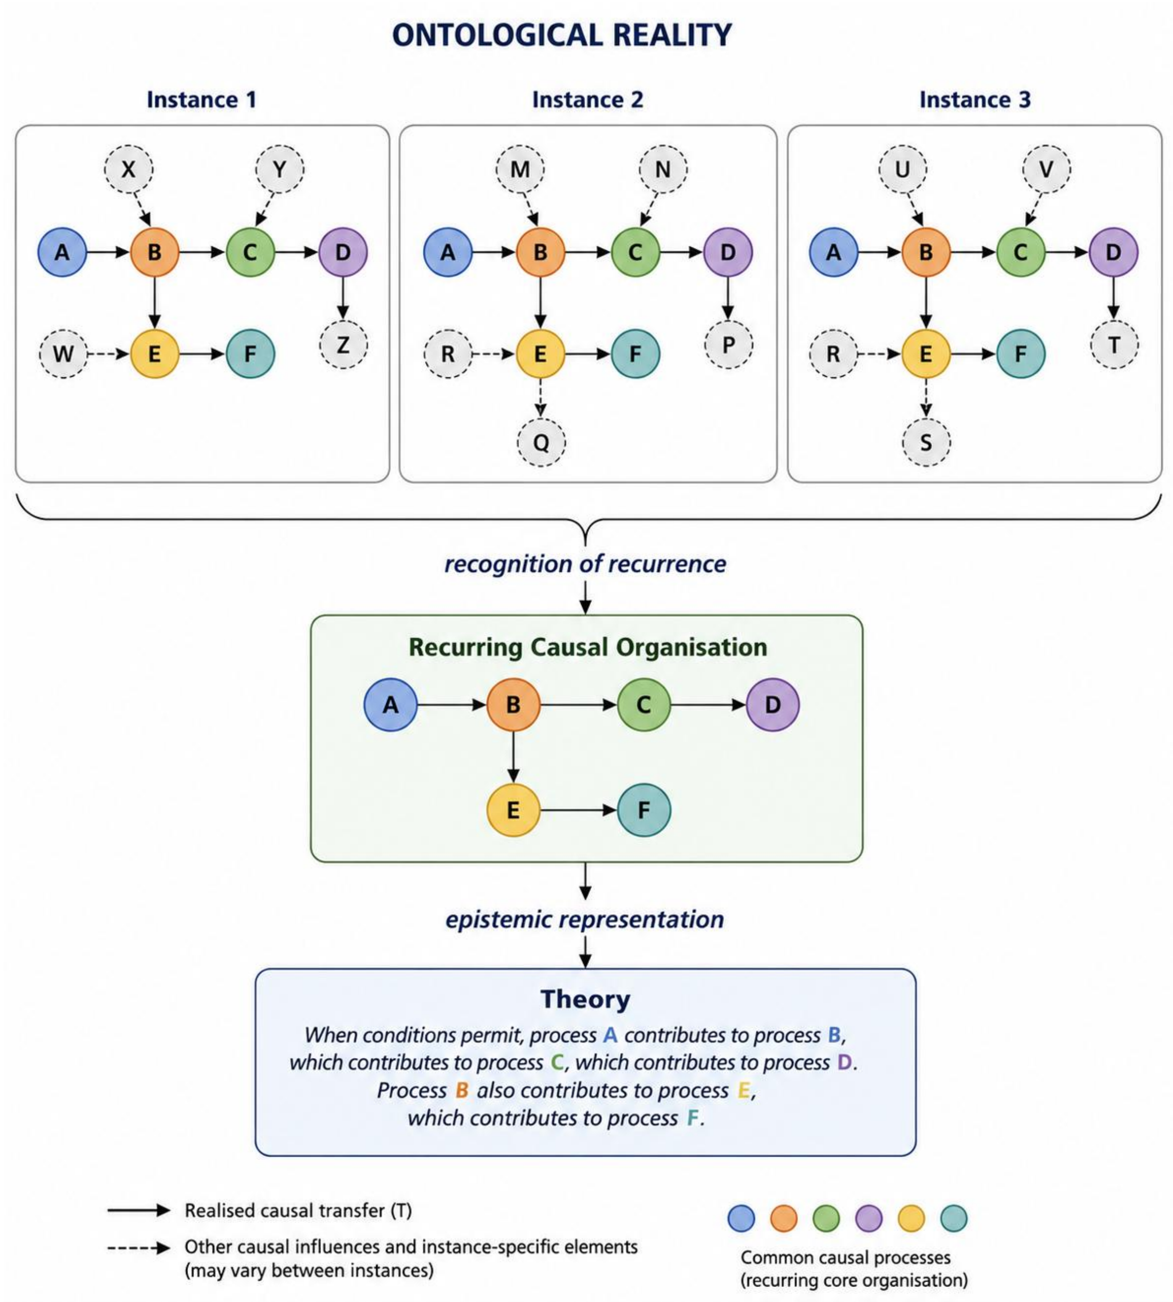

Figure 12. From Causal Reality to Theory.
Similar causal organisation may recur within otherwise different instances of reality. Recognition of this ontological recurrence allows the common causal organisation to be represented epistemically as theory. A theory therefore need not reproduce the complete causal organisation of individual instances, but may represent those causal relationships that recur across them.

11.4 Epistemic and Ontological Causality

Hypotheses, explanations and theories do not exist in epistemic isolation. Once represented within human cognition or other information-processing systems, they may interact with other beliefs, observations and representations. One theory may contribute to the formation of another, an observation may cause a hypothesis to be revised, and two previously separate explanations may be combined to produce a new explanation.

There are therefore causal relationships within the epistemic domain. Hypotheses, explanations, theories and other beliefs form changing networks in which existing representations contribute causally to the formation, modification or rejection of others. These epistemic causal relationships are themselves physically instantiated through the underlying processes and transfers occurring within the agents or systems in which the representations exist.

Epistemic representations may also contribute to effects beyond the epistemic domain. A belief about the likely consequences of an action may influence an agent's decision, resulting in neural activity that controls muscular activity and ultimately produces causal transfers into the external environment. Through the causal leverage described in Section 10, relatively small informationally structured processes within an agent may therefore influence much larger external flows of matter and energy.

A theory can consequently become ontologically causal through the agency of those who act upon it. An engineering theory may contribute to the construction of a bridge; an economic theory may influence government policy; and a scientific theory may guide an experiment that physically changes the world. The theory does not act independently of its physical instantiation and the agents involved, but it forms part of the causal network through which those effects arise.

The relationship between causal reality and causal knowledge is therefore recursive. Causal reality gives rise to observations and epistemic representations; those representations interact within the epistemic domain and, through the causal capabilities of agents, may subsequently contribute to changes in causal reality.

12. Implications for General Systems Theory

The causal framework developed in this paper provides a common foundation for many of the concepts introduced in the author's previous paper, "Foundations of General Systems Theory from a Critical Realist Perspective" (Challoner, 2026). Entities and their properties exist in configurations; processes produce change; and transfers of matter, energy and information physically connect processes into causal networks. Extended causal reality can therefore be understood as an alternating organisation of processes and transfers:

… P → T → P → T → P → T → P …

As illustrated in Figure 7, PTP and TPT provide complementary epistemic perspectives upon this same underlying causal reality. PTP foregrounds interactions between processes through transfers, while TPT foregrounds processes transforming inputs into outputs. Systems and their interactions can therefore be represented within a common causal framework without requiring different forms of causality.

Circumstances, conditions and constraints can also be incorporated into this framework. Rather than being regarded simply as passive background features, they may arise from persistent causal organisations. Conditions enable particular causal processes, while constraints restrict or inhibit processes or restrict the pathways through which causal transfers can occur. Persistent configurations similarly maintain possibilities for future causal interaction.

Systems can consequently be understood as relatively coherent causal organisations within the wider causal networks of reality. Persistent configuration establishes continuing possibilities for interaction; processes and transfers realise some of those possibilities; and differential coupling produces concentrations of causal activity. Where these organised interactions generate emergent properties, behaviours or outputs attributable to the whole, the resulting causal organisation may be recognised as a system or subsystem.

Emergence can in turn create new causal capabilities. A system may generate transfers and effects that its components could not produce independently, while the system itself may subsequently become a component of a still larger causal organisation. This produces an ontological hierarchy of causal organisation without requiring any departure from physical causality: higher-level causal capabilities remain instantiated through the organised processes and transfers of their components.

Causal leverage further explains how relatively small causal inputs can control much larger independently available flows. This provides a physical foundation for agency: informationally structured neural activity can control muscular activity, which can alter external conditions or constraints and thereby enable, inhibit or redirect much larger flows of matter and energy.

The framework also connects causal ontology with causal knowledge. Recurring relationships within causal reality may be recognised and represented as hypotheses, explanations and theories. These representations can themselves interact causally within the epistemic domain and, through the leverage capabilities of agents, contribute to subsequent changes in causal reality.

Causal organisation therefore provides a bridge between the foundational ontology established in “Foundations of General Systems Theory from a Critical Realist Perspective” and subsequent developments of the framework. It provides foundations for the treatment, in later papers, of emergence, cognition and agency, hypothesis testing, and the diagnosis and intervention of complex systems. Figures 7 and 11 summarise particularly important aspects of this argument: the underlying continuity of causal reality and the capacity of organised causal systems to leverage relatively small inputs into much larger effects.

In this way, causality is not an additional concept attached to systems theory. It is the physical continuity through which configurations become dynamic, processes interact, systems operate, emergent capabilities arise, agents act, and changes propagate through reality.

13. Conclusion

The central conclusion of this paper is that systems theory and causal theory are strongly interrelated. A system cannot be fully understood without considering the causal organisation through which its components interact, while causal relationships in complex reality cannot readily be understood without considering the systems and persistent causal organisations through which they occur.

Causality is physically instantiated. Processes are connected through transfers of matter, energy and information, and these interactions concatenate into the causal networks through which change propagates. Information does not require a separate non-physical form of causation because informational structure is embodied in matter or energy and becomes causal when processes respond to that structure.

Systems provide relatively persistent concentrations of this causal organisation. Their configurations establish causal possibilities, their processes and transfers realise those possibilities, and their organisation can generate emergent causal capabilities unavailable to their components independently. Systems can themselves become components of larger systems, producing an ontological hierarchy in which new causal capabilities arise at successive levels of organisation.

The distinction between causal reality and causal knowledge is equally important. Causes, effects, systems, events and causal boundaries are partly ways in which observers select and represent portions of a continuous causal reality. PTP and TPT demonstrate this particularly clearly: they are alternative epistemic perspectives upon the same underlying sequence of processes and transfers. Hypotheses, explanations and theories similarly represent selected and recurring causal organisations rather than constituting causal reality itself.

Yet ontology and epistemology are not causally separate. Epistemic representations are physically instantiated within agents, interact with other beliefs and representations, and may influence subsequent action. Through causal leverage, ideas can therefore contribute to changes in the physical world without requiring any departure from physical causality.

The relationship between systems and causality is consequently reciprocal. Causal organisation helps explain what systems are and how they operate; systems provide persistent organisations through which causality is realised, maintained and amplified. Bringing the two together provides a common foundation upon which subsequent treatments of emergence, cognition, agency, investigation and intervention can be developed.

References

Bhaskar, R. (1978). *A realist theory of science* (2nd ed.). Harvester Press.

Challoner, J. A. (2026). *Foundations of systems theory from a critical realist perspective*. Rational Understanding.

Dowe, P. (2000). *Physical causation*. Cambridge University Press.

Hitchcock, C. (2018). Probabilistic causation. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2018 ed.). Metaphysics Research Lab, Stanford University.

Juarrero, A. (1999). *Dynamics in action: Intentional behavior as a complex system*. MIT Press.

Landauer, R. (1991). Information is physical. *Physics Today, 44*(5), 23–29. https://doi.org/10.1063/1.881299

Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science, 67*(1), 1–25. https://doi.org/10.1086/392759

Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Clarendon Press.

Motter, A. E., & Lai, Y.-C. (2002). Cascade-based attacks on complex networks. *Physical Review E, 66*(6), 065102. https://doi.org/10.1103/PhysRevE.66.065102

Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). Cambridge University Press.

Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson.

Rothman, K. J. (1976). Causes. *American Journal of Epidemiology, 104*(6), 587–592.

Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world*. Princeton University Press.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society, 106*(6), 467–482.

von Bertalanffy, L. (1968). *General system theory: Foundations, development, applications*. George Braziller.

Wiener, N. (1948). *Cybernetics: Or control and communication in the animal and the machine*. John Wiley & Sons.

Definitions

D6.1 – Cause and Effect
Cause and effect are epistemic roles attributed to causal phenomena when portions of causal reality are represented for explanatory purposes.
The same causal phenomenon may therefore be represented as an effect in relation to preceding causality and as a cause in relation to subsequent causality.

D5.5 – Interaction
Participation by two or more entities in a causal relationship involving the transfer of matter, energy, or information between them.

D5.7 – Transfer
The physical passage or transmission of matter, energy, or information through a causal relationship.

D6.2 – Necessary Cause
A necessary cause is a causal contribution without which a specified process or effect cannot occur.

D6.3 – Sufficient Cause
A sufficient cause is a causal contribution, or combination of causal contributions, that provides everything required for a specified process or effect to occur in the absence of sufficient inhibition.

D6.4 – Enabling Cause or Enabler
An enabling cause is a causal contribution that allows a process to operate or an effect to occur.

D6.5 – Inhibiting Cause or Inhibitor
An inhibiting cause is a causal contribution that prevents, restricts or disrupts a process that would otherwise be capable of operating.

D6.6 – Causal Chain
A causal chain is a sequence of causally connected causes and effects in which an effect may subsequently become a cause of a further effect.

D2.23 – Network

A group of causally interconnected entities or relationships.

D6.7 – Causal Network
A causal network is an interconnected causal organisation of processes and transfers in which causal pathways may branch, converge and interact. It is a physically connected organisation extending through space and time, rather than merely a logical arrangement of causes and effects.

D6.8 – Direct Causation
Direct causation occurs when a cause influences an effect through an identifiable causal interaction without an intervening causal system or event.

D6.9 – Indirect Causation
Indirect causation occurs when causal influence passes through one or more intermediate processes or systems before producing the effect under consideration.

D6.10 – Process–Transfer–Process (PTP) or Event-based Perspective
PTP is an epistemic representation that places a conceptual boundary around two processes and the causal transfer connecting them, thereby foregrounding causal interaction, transfer and change.

D6.11 – Transfer–Process–Transfer (TPT) or Systems-based Perspective
TPT is an epistemic representation that places a conceptual boundary around a process together with its incoming and outgoing transfers, thereby foregrounding the transformation of inputs into outputs.

D2.14 – Structure (Technical) (Updated)
The overall pattern exhibited by the causal organisation of an assembly or system.

D6.12 – Causal Structure
An epistemic representation of a recurring pattern of causal organisation recognised within causal networks.

D6.13 – Linear Causality
Linear causality is a causal structure in which a cause produces an effect through a single causal pathway.

D6.14 – Causal Tree
A causal tree is a causal structure in which one cause gives rise to several effects through branching causal pathways.

D6.15 – Causal Root
A causal root is a causal structure in which several causal pathways converge upon a single effect.

D6.16 – Domino Causality
Domino causality is a sequential causal structure in which each effect becomes the cause of the next effect in the sequence.

D6.17 – Cyclic Causality
Cyclic causality occurs when a sequence of causal relationships repeatedly recreates the same types of causes and effects.

D6.18 – Mutual Causality
Mutual causality occurs when two entities or processes causally influence one another.

D6.19 – Feedback
Feedback occurs when an output from a process or system returns, directly or indirectly, as an input that influences the subsequent operation of that process or system.

D6.20 – Cascading Causality
Cascading causality is a causal chain in which each successive stage produces both the next stage in the chain and an additional contribution to a common accumulating effect.

D6.21 – Propagating Causality
Propagating causality occurs when causal effects spread through a connected network, producing further effects that initiate additional causal pathways elsewhere in the network.

D6.22 – Relational Causality
Relational causality occurs when an effect depends upon a relationship or difference between the states of two or more entities rather than upon the state of either considered alone.

D6.23 – Threshold Causality
Threshold causality occurs when an effect is produced only when a relevant variable reaches or crosses a particular value.

D6.24 – Transient Causal Network
A transient causal network is a causal organisation associated with a particular causal episode that develops and subsequently largely ceases to operate.

D6.25 – Persistent Causal Network
A persistent causal network is a causal organisation that continues to operate over an extended period despite changes in its constituent matter, energy, information, processes or components.

D3.9 – Constraint or Condition
Constraints and conditions are complex causal networks delivering multiple enabling or inhibiting inputs to a system, thereby defining the range of behaviours in which the system can or cannot engage. The precise causal mechanisms involved are often unknown, distributed, or analytically impractical to specify directly, requiring simplification through the concepts of constraint and condition. These do not denote fundamentally different ontological categories, but different emphases within organised causal systems. In the case of constraints, the emphasis lies on the range of behaviours excluded or limited. In the case of conditions, it lies on the range of behaviours enabled or permitted.

D6.26 – Circumstance

A feature of the causal context within which an event or process occurs, which may function as a condition or constraint in relation to that event or process.

D6.27 – Differential Coupling
Differential coupling occurs where causal interactions are more frequent, stronger or faster among some processes than between those processes and others within the wider causal network.

D6.28 – Causal Leverage
Causal leverage occurs where a relatively small causal input of matter, energy or information modifies a process or constraint that regulates a larger independently available causal flow. Leverage may operate through enablement, inhibition or control.

D6.29 – Causal Mechanism
The causal organisation of processes, transfers, conditions and constraints through which a particular effect or behaviour is produced, sustained, modified or inhibited.

D6.30 – Causal Hypothesis
A causal hypothesis is a provisional epistemic representation of possible causal organisation that may account for an observed or anticipated phenomenon and can be tested against causal reality.

D6.31 – Causal Explanation
A causal explanation is an epistemic representation of causal organisation considered sufficiently well supported to account for how or why a phenomenon occurs.

D6.32 – Causal Theory
A causal theory is an epistemic representation of causal organisation believed to recur across a class of phenomena.

Propositions

P6.1 – Physical Instantiation of Causality
Causal influence between separate processes is physically instantiated through transfers of matter, energy or physically embodied information. No causal influence can pass between separate processes without a physical transfer connecting them.

P6.2 – Causal Continuity
Causal reality forms continuing causal networks in which effects may become subsequent causes; individual cause-and-effect relationships are therefore portions of a wider causal organisation rather than necessarily independent causal units.

P6.3 – Transfer-Level Necessity
Necessity and sufficiency attach most fundamentally to the causal transfers required by recipient processes: a required transfer remains necessary irrespective of the source from which it originates.

P6.4 – Alternative Source Proposition
Where alternative source processes can provide equivalent necessary transfers, no individual source need be necessary even though the transfer itself remains necessary.

P6.5 – Sufficiency of Complete Enabling Inputs
Where the complete set of necessary enabling transfers is present and no sufficient inhibitor prevents operation, their combination is sufficient for the recipient process to occur and produce the specified effect.

P6.6 – Causal Precedence of Sufficient Inhibition
Sufficient inhibition takes causal precedence over sufficient enablement: a sufficiently inhibiting causal contribution can prevent or limit a process even when its complete set of enabling inputs is present.

P6.7 – Epistemic Probabilism Proposition
An observed probabilistic causal relationship may arise from incomplete knowledge of other necessary enabling or inhibiting causal contributions and therefore does not necessarily imply intrinsically probabilistic underlying causation.

P6.8 – Intervention Opportunity Proposition
As the number of intermediate causal interactions in an indirect causal pathway increases, the number of opportunities for that pathway to be enabled, inhibited, redirected or interrupted also increases.

P6.9 – Causal Representation Proposition
The underlying causal organisation exists independently of its representation, while the causes, effects, boundaries and level of causal detail represented depend partly upon the purpose and granularity of observation.

P6.10 – Alternating Causal Organisation
Extended causal reality can be represented ontologically as an alternating organisation of processes and transfers: ... **P → T → P → T → P** ....

P6.11 – Structural Equivalence of Cause and Effect
Because the same causal phenomenon may occupy the relational role of effect with respect to preceding causality and cause with respect to subsequent causality, causes and effects within a continuing causal representation should possess compatible causal structure.

P2.9 – Dual Representation of Causality Proposition

Causal processes may be represented either as transfers between systems (process-transfer-process, PTP) or as transformations within systems involving inputs and outputs (transfer-process-transfer, TPT). These are complementary representations of the same causal structure.

P6.12 – Shared-Component Causal Continuity
Successive PTP cause-and-effect units maintain causal continuity by sharing a process, whereas successive TPT cause-and-effect units maintain continuity by sharing a transfer.

P6.13 – PTP Necessity and Sufficiency Proposition
PTP provides the most direct representation of necessary and sufficient causation because causal pathways converge upon recipient processes, allowing required transfers and their alternative sources to be distinguished.

P6.14 – PTP–TPT Systems–Causality Bridge
PTP provides an event-based perspective upon causal interaction, while TPT provides a systems-based perspective upon process functioning; together they provide complementary representations that bridge causal theory and systems theory.

P6.15 – Causal Cognition Hypothesis
TPT may facilitate recognition of recurring patterns in process behaviour, whereas PTP may facilitate explicit logical analysis and decompression of the causal networks through which those behaviours arise.

P6.16 – Persistent Causal Context Proposition
Persistent causal networks provide part of the causal context within which transient events occur and may establish circumstances, conditions, constraints and continuing causal possibilities.

P6.17 – Configuration–Realisation Proposition
Configuration establishes potential causal pathways, while processes and transfers realise particular pathways within those possibilities.

P3.5 – Organised Configuration Proposition
Organised configurations arise where causal interactions between entities produce recurring or stable configurational arrangements.

P6.18 – Differential Coupling and System Recognition
Differential coupling can produce relatively coherent concentrations of causal organisation that may be recognised as systems or subsystems despite continuing causal interaction across their boundaries.

P6.19 – Emergence Proposition
Emergent properties arise from the organisation and causal interaction of components and are not exhibited by those components considered independently.

P5.6 – Emergent Systems
A system emerges when organised causal interactions produce properties, behaviours or outputs that are not exhibited by its interacting components in isolation.

P5.7 – Emergent Causation
Emergence gives rise to new causal capabilities that are not available to the interacting components considered in isolation.

P6.20 – Emergent Entity Transfer Proposition
Entities produced through the organised causal interactions of a source system may themselves be transferred as organised wholes and subsequently participate in further causal interactions.

P5.8 – Recursive Organisation
Systems are recursively organised such that they may comprise interacting subsystems while simultaneously functioning as components of larger systems.

P6.21 – Causal Leverage Proposition
The causal significance of an input depends upon its location and effect within causal organisation rather than simply upon its physical magnitude; relatively small inputs may therefore produce disproportionately extensive consequences by enabling, inhibiting, controlling or redirecting larger causal flows.

P6.22 – Physical Basis of Agency
Agency can be understood physically as the capacity of an organised system to generate causal transfers that selectively influence subsequent causal pathways.

P6.23 – Causal Decompression Proposition
Causal investigation can progressively decompress recognised associations by identifying and testing the processes, transfers, conditions and constraints through which the represented phenomena are causally connected.

P6.24 – Epistemic Representation Proposition
Hypotheses, explanations and theories are epistemic representations of selected causal organisation rather than causal reality itself; they differ principally in epistemic status, purpose and scope.

P6.25 – Partial Revision Proposition
Because causal hypotheses, explanations and theories represent decomposable causal organisation, evidence may support some components while contradicting others, allowing representations to be progressively revised rather than accepted or rejected only as indivisible wholes.

P6.26 – Recurrence and Explanatory Power
The explanatory reach of a theory depends partly upon the extent to which the causal organisation it represents recurs across instances of reality.

P6.27 – Cross-Domain Explanatory Transfer
Where sufficiently similar causal organisation recurs in different types of system, a theory developed in one domain may contribute to explaining phenomena in another.

P6.28 – Recursive Ontology–Epistemology Proposition
Causal reality gives rise to epistemic representations which, when physically instantiated within agents or other information-processing systems, may themselves participate in causal networks and contribute to subsequent changes in causal reality.

P6.29 – Systems–Causality Reciprocity
Causal organisation helps explain what systems are and how they operate, while systems provide persistent organisations through which causality is realised, maintained and amplified.

