




ELSEVIER

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Simulation Modelling Practice and Theory

journal homepage: www.elsevier.com/locate/simpat

Structural Analysis of System Dynamics Models

Lukas Schoenenberger^{a,*}, Alexander Schmid^b, Radu Tanase^c, Mathias Beck^d,
Markus Schwaninger^e

^a *Institute of Health Economics and Health Policy, Bern University of Applied Sciences, Murtenstrasse 10, 3008 Bern, Switzerland*

^b *Institute Digital Enabling, Bern University of Applied Sciences, Brückenstrasse 73, 3005 Bern, Switzerland*

^c *URPP Social Networks, University of Zurich, Andreasstrasse 15, 8050 Zurich, Switzerland*

^d *Department for Management, Technology, and Economics, KOF Swiss Economic Institute, ETH Zurich, Leonhardstrasse 21, 8092 Zurich, Switzerland*

^e *Institute of Strategy and Management, University of St. Gallen, Dufourstrasse 40a, 9000 St. Gallen, Switzerland*

ARTICLE INFO

Keywords:

System dynamics
structural analysis methods (SAM)
graph theory
network science
network controllability
path analysis
algorithmic detection of archetypal structures (ADAS)
policy design

ABSTRACT

System dynamics (SD) is an established discipline to model and simulate complex dynamic systems. The primary goal of SD is to evaluate and design new policies that can impact the system under study in a desired way. Policy design, that is, identifying effective model levers, however, is a challenge and in many cases trial-and-error driven. In this article, we introduce a new and coherent framework for model analysis, called structural analysis methods (SAM), to facilitate the policy design process in complex SD models. SAM provides a resource-efficient and effective means for the detection of candidate policy parameters. It enables to identify intended and unintended effects of activating these policy parameters, and to discover candidate structural changes such as introducing new variables and links in SD models. The main innovation of SAM is that it translates the structure of SD models into weighted digraphs allowing algorithmic tools from the realms of graph theory and network science to be applied to SD. SAM is validated on the basis of two well-known simulation models of increasing complexity: the third-order Phosphorus Loops in Soil and Sediment (PLUM) model and the fifth-order World2 model. The validation shows that SAM seems to be most valuable for the analysis of more complex simulation models (World2) and is less suited for the analysis of low complexity models (PLUM).

1. Introduction

Recently, system dynamics (SD)—an approach to modelling and simulating complex systems—alongside with other systems thinking methods has gained remarkable attention in academia and practice, namely in the fields of health care [1–3] and sustainability [4–6]. This might be explained by the distinctive advantages SD offers to users that seem to be more appreciated nowadays than in the past. Among others, key advantages of SD are (i) an accessible graphical iconography and user-friendly modelling software, (ii) the capacity to engage stakeholders throughout the modelling process, and (iii) the development of aggregated causal (simulation) models that are particularly useful for policy analysis and design [7–9].

In SD, model building and analysis usually entails some or all of these standard activities (see Fig. 1) [9, 7]. In the beginning, the problem needs to be articulated clearly in terms of theme selection, key variables, time horizon to be considered, and historical

* Corresponding author.

E-mail address: lukas.schoenenberger@bfh.ch (L. Schoenenberger).

<https://doi.org/10.1016/j.simpat.2021.102333>

Received 5 November 2020; Received in revised form 8 April 2021; Accepted 9 April 2021

Available online 15 April 2021

1569-190X/© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

behaviour of the key concepts and variables. In the next phase a so-called dynamic hypothesis is formulated that represents modelers' understanding about the key causal mechanisms that generate the problem at hand. Typically, the dynamic hypothesis is expressed using some form of mapping technique such as causal loop diagrams (CLDs), stock-and-flow diagrams (SFDs), model boundary diagrams etc. Problem articulation and the formulation of a dynamic hypothesis belong to the qualitative part of SD. Next a simulation model is developed in order to test the explanatory power of the dynamic hypothesis in conjunction with the problem at hand. For this purpose, the simulation model needs to pass a couple of standard validation checks to build up trust into the usefulness of the model. Finally, the validated simulation model is used for policy design and evaluation. The last three phases of simulation model formulation, validation, and policy design, belong to the quantitative part of SD.

In this context, policy design refers to the process of finding influential parameters and model relationships, as well as identifying possible structural changes (i.e., addition or elimination of parameters, variables or relationships) in SD models which can—if manipulated or implemented cleverly—steer the behaviour of key outcome variables into a desired direction. The identification of influential parameters or model levers, especially in large (complex) SD models, however, is a challenge. This has to do with the high degree of interconnected model variables and non-linear relationships (links) typically present in SD models. Consequently, in many cases, the policy design process in SD is far from trivial and still trial-and-error driven [10–12].

In this article, we introduce the *structural analysis methods* (SAM)—a coherent model analysis framework which integrates powerful analysis tools that have been developed previously—for the purpose of facilitating the policy design process in SD. SAM is centred on the examination of model *structure* and provides a resource-efficient (in terms of the amount of required information/computational complexity) and effective means for policy design in complex models. The main assumption of structural analysis is that the structure of SD models can be accurately described through directed weighted networks (weighted digraphs), making it accessible to algorithmic exploration by means of network science and graph theory [13–15, 12]. This implies that variables and causal relationships in SD models can be converted into vertices (nodes) connected by edges (links). Beyond SD models, the same assumption can also be made for cognitive [16] and causal maps [17] increasing the potential reservoir of models where SAM might be applied. To the best of our knowledge, SAM is the first attempt to combine SD with network science—two important fields in the sphere of modelling and analysing complex systems.

The remainder of the paper is structured as follows: the next section briefly introduces the basic concepts of SD. Then, the methodological and graph-theoretical background of SAM are presented in sections three and four. Section five introduces the analysis tools—network controllability, path analysis, and algorithmic detection of archetypal structures (ADAS)—in detail. Section six validates SAM on the basis of two established SD models: Phosphorus Loops in Soil and Sediment (PLUM) [18] and World2 [19]. Limitations of SAM are presented in section seven. The last section discusses key findings and provides conclusions.

2. Basic Concepts of System Dynamics (SD)

SD was created by Jay W. Forrester at the Massachusetts Institute of Technology during the mid-1950s.¹ His key idea was that any complex system can be modelled by no more than two kinds of variables²: stocks and flows [20, 7]. A stock is an accumulation of units and graphically represented as a rectangle³, a flow, which can be an inflow or an outflow, implies a movement of units per time and is visually illustrated by an arrow (see Fig. 2). Stocks can only be changed through their connected in- and outflows. Flows in turn are regulated by decision functions, visualized as valves on in- and outflows.

Simulations are in continuous time. The corresponding integral equation is:

$$Stock(t) = \int_{t_0}^t [Inflow(s) - Outflow(s)] ds + Stock(t_0) \quad (1)$$

Using the bathtub metaphor, the valve would regulate the flow of water. The shapes at the edges of the diagram symbolise a source and a sink. Both represent areas outside the model, – black boxes, which are not analysed. One might at some point dissolve part of a source or sink: for example, if the model represents an industrial firm, an extension of the model could represent additional modules along the value chain, – backward or forward.

Forrester showed that complex systems are largely composed by causal loops, and that they show delays (see Fig. 3). They are closed in that they are influenced by their own past behaviours [21]. Hence, different from other sciences, including most of economics, the perspective of SD is one of endogenous causation [22].

A system model usually consists of several such circuits that are interconnected. The modelled structure of the system determines its simulated behaviour. Due to the closed loops and the delays, nonlinear behaviour patterns result, which may be counterintuitive in nature.

SD offers a formal apparatus rooted in systems theory. It embodies feedback thought, and it builds on the traditions of control theory and servomechanism theory [23]. SD offers the language of a formal science. It is generic and broadly applicable, and better suited for dealing with complex systems, than pure mathematics and formal logic. Ideal ambits of application for SD are social,

¹ Jay W. Forrester (1918–2016) was a pioneering systems scientist and computer engineer, professor and founder of the Systems Dynamics Group, at the Massachusetts Institute of Technology.

² Auxiliary variables and parameters not included.

³ Typical stocks are resources or assets (material, immaterial), goods, inventories. Examples: natural resources, financial assets, personnel, experience, knowledge, etc.

Phase 1	Phase 2	Phase 3	Phase 4	Phase 5
Problem articulation	Dynamic hypothesis formulation	Simulation model development	Simulation model validation	Policy design and evaluation
<i>Qualitative SD</i>			<i>Quantitative SD</i>	

Fig. 1. SD model building and analysis process.



Fig. 2. Minimal stock-and-flow diagram (SFD).

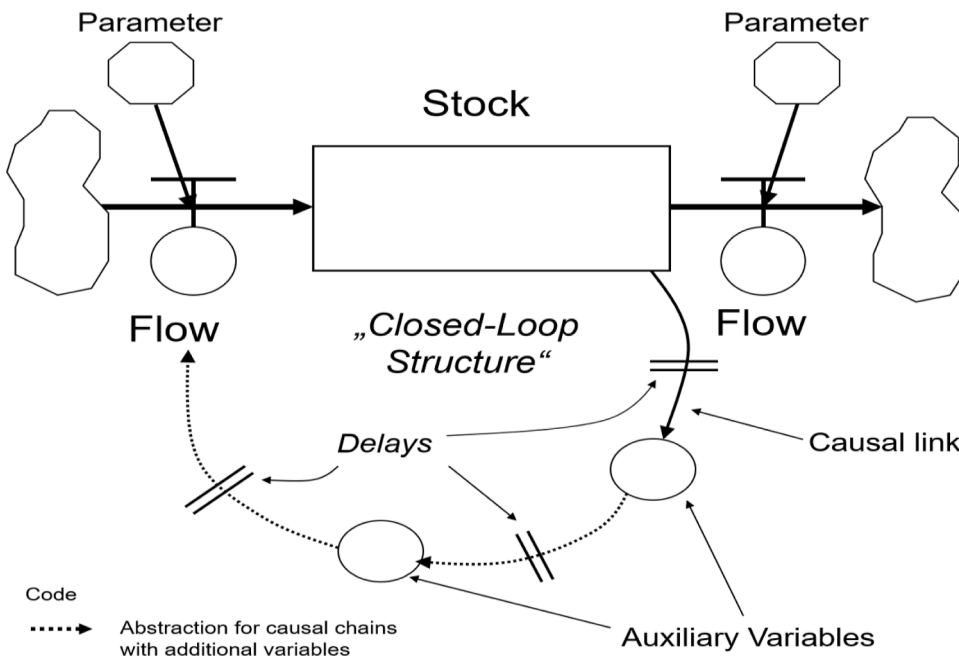


Fig. 3. Closed-loop stock-and-flow diagram (SFD).

socio-technical, and ecological systems.

3. Methodological Background

SAM consists of three distinct but interrelated tools that are best applied in the following sequence in order to support the policy design process in SD:

1. Detection of candidate policy parameters (model levers) in SD models based on network controllability calculations.
2. Detection of intended and unintended consequences of triggering candidate policy parameters in SD models using the path analysis method.
3. Detection of candidate structural changes (policy links) in SD models by means of the algorithmic detection of archetypal structures (ADAS) method.

In the first step, SD models are scanned for candidate policy parameters (model levers) using network controllability, an emerging field in network science and first described by Liu, Slotine and Barabasi [24] in *Nature*. To that end, we calculate the *control centrality*, that is, a measure indicating the ability of a single node in controlling an entire network, for every variable in an SD model of interest [25]. The feasibility and usefulness of deriving the control centrality for variables in SD models was recently shown by Schoenenberger and Tanase [12] analysing one prominent case study model. In this case study, they showed that variables with a high control centrality can exert a more substantial impact on model behaviour of key outcomes than the ones with a low control centrality.

In the second step, intended and unintended consequences of changing selected candidate policy parameters on outcome variables are identified. This can be achieved through path analysis for which algorithmic methods derived from graph theory have been formulated [26, 27]. The value of path analysis for causal inference in qualitative SD has been demonstrated in diverse modelling contexts such as health care [28] and counter-terrorism [29].

In the third and final step, an SD model is examined for the existence of archetypal structures using the ADAS method [15, 30]. Archetypal structures are generic structures that appear in any system regardless of its size and context. Typically, from the intervenor's perspective, archetypal structures cause problematic system behaviour because they embody structural mechanisms that have the tendency to nullify or even revert the impact of well-intended system interventions (policies). These archetypal structures have been described in SD notation using variables and causal relationships [31, 32]. Consequently, it is possible to identify these (sub-)structures as part of larger SD models. The ADAS method builds on two pattern recognition algorithms that can detect the four archetypal structures as defined by Wolstenholme [31] in SD models. This makes ADAS a tool that is valuable for both system diagnosis and policy design. It supports system diagnosis because it can direct modelers' attention towards archetypal structures as possible explanations for unexpected system behaviour; and it can also facilitate the design of effective policies through so-called solution archetypes. Solution archetypes propose specific structural changes (i.e., addition of new links, variables, loops etc.) to SD models that compensate for the negative consequences associated with archetypal structures.

We conceive of SAM as a basket of formal analysis methods that complement well-established tools in SD such as pathway participation metric (PPM) [36–38], model structure analysis (MSA) [13], or eigenvalue elasticity analysis (EEA) [33–35]. Fig. 4 illustrates how SAM fits into the larger scheme of formal analysis methods in SD. Crucially, while PPM and EEA link model structure to model behaviour, SAM and MSA focus on model structure exclusively. Consequently, SAM and MSA enable a less fine-grained model analysis compared to the other two methods but they are superior in case of qualitative model analysis where behavioural information is missing. An introduction into PPM, MSA, and EEA is beyond the scope of this article, but we refer the interested reader to the sources indicated in Fig. 4.

4. Graph-theoretical Background

4.1. Capturing Variables and Relationships in Adjacency Matrices

SAM builds on the assumption that the structure of an SD model can be precisely characterized as a directed weighted network or weighted digraph G consisting of a set of vertices (nodes) V and a set of edges (links) E [13–15]. Every edge e_{ij} linking two vertices $v_i, v_j \in V$, that is, $v_i \rightarrow v_j = e_{ij} \in E$, denotes a direct causal relationship. V and E can be captured in an adjacency matrix A , a $|V| \times |V|$ square matrix with $A = (e_{ij})$ [39]. Each vertex appears twice, both in a row and a column of A . The values in row A_i represent the successor set for vertex v_i , while the values in column A_i represent the predecessor set for v_i . Put simply, A captures information with respect to the number of variables and their interrelations in an SD model. Fig. 5 shows an example of a stock-and-flow diagram (SFD), its corresponding graph representation and adjacency matrix.

Compared with the SFD, the digraph only specifies causal relationships (arrows), without distinguishing between stocks and flows. The adjacency matrix then translates these arrows (“causes”) into binary digits: reading row-wise, “1” denotes an outgoing arrow, “0” no outgoing arrow. The cells in the matrix represent edges, which specify causalities.

4.2. Capturing Polarity of Relationships and Delays as Edge Weights

The information captured in the adjacency matrix A is sufficient for the algorithmic detection of candidate policy parameters in SD models. However, in order to algorithmically detect both intended and unintended consequences of policy parameters and archetypal structures (structural changes), information beyond the connectivity of variables provided in A is needed. Specifically, information is required about the polarity of relationships as well as the existence and magnitude of delays. To account for the polarity of relationships, the following edge weights p_{ij} are used [15]:

$$p_{ij} = \begin{cases} -1 & \text{if } v_i \rightarrow v_j \\ 1 & \text{if } v_i \overset{+}{\rightarrow} v_j \end{cases} \quad (2)$$

$v_i \rightarrow v_j$ implies that, *ceteris paribus*, an increase in the cause variable v_i causes a decrease in the effect variable v_j *below what it would otherwise have been* or vice versa, so the change is in the opposite direction, corresponding to a “negative” relationship (negative link polarity) in SD models. In contrast, $v_i \overset{+}{\rightarrow} v_j$ indicates that, *ceteris paribus*, an increase in the cause variable v_i induces an increase in the effect variable v_j *above what it would otherwise have been* or vice versa, so the change is in the same direction, which corresponds to a “positive” relationship (positive link polarity) in SD models.

Relationships in SD models are also characterized by the presence and magnitude of delays [7]. To include temporal information in




Fig. 4. Integration of the structural analysis methods (SAM) into formal model analysis in system dynamics (SD) (adapted from [12]).

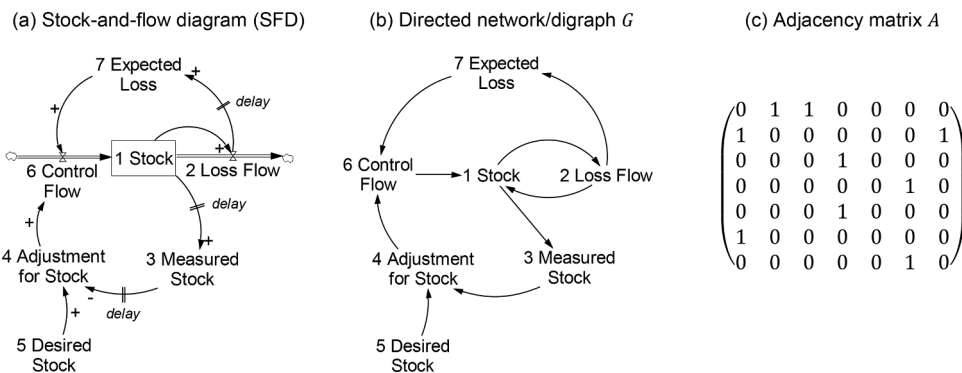


Fig. 5. Stock-and-flow diagram (SFD) with corresponding directed network/digraph representation and adjacency matrix.

a weighted directed network/digraph, an additional set of edge weights τ_{ij} is used [15]:

$$\tau_{ij} = \begin{cases} 1 & \text{if a change in } v_i \text{ immediately affects } v_j \\ 2 & \text{if } v_i \text{ is a stock variable} \\ 4 & \text{if a change in } v_i \text{ affects } v_j \text{ only after a significant delay} \end{cases} \quad (3)$$

whereas τ_{ij} indicates the time delay between a change in the cause variable v_i and its impact on the effect variable v_j . If a change in the cause variable v_i immediately impacts the effect variable v_j , then $\tau_{ij} = 1$. If the cause variable v_i is a stock variable, then $\tau_{ij} = 2$, meaning the effect variable v_j reacts with a slight delay to any change in v_i . Finally, if a change in the cause variable v_i impacts the effect variable v_j only after a significant time delay, then $\tau_{ij} = 4$. Values corresponding to both p_{ij} and τ_{ij} are stored in two separated weighted adjacency matrices named *polarity matrix P* and *temporality matrix T* [30]. The distinction between positive and negative polarities is made in the *polarity matrix P*, but not in the *temporality matrix T*, where only the strength of delays counts. Fig. 6 presents the same example SFD as in Fig. 5 together with the corresponding polarity and temporality matrices.

The value set associated with τ_{ij} requires more explanation. The values τ_{ij} seem rather arbitrary when interpreted in an absolute sense and could be exchanged by other numbers. However, what is relevant is the ratio between the values and how they relate to SD practice. In SD, causal relationships that operate instantaneously are probably the most common case. Because the effect is immediate, the delay is minimal and thus is specified as $\tau_{ij} = 1$. Then, SD also offers mathematical functions to explicitly account for delays if they are pronounced and important to the modelled system. Schoenenberger, Schmid and Schwaninger [15] proposed to code these functions with $\tau_{ij} = 4$, simply indicating that the reaction time of the effect variable v_j is four times greater than in the instantaneous case. Finally, stock variables, i.e., accumulations, are of central importance to the SD approach. Typically, stocks accumulate or deplete gradually and so introduce memory and inertia into (modelled) systems [7]. To account for this fact, in case the cause variable v_i is a stock, Schoenenberger, Schmid and Schwaninger [15] set $\tau_{ij} = 2$, positioning this delay between the first two.⁴

Overall, the set of edge weights τ_{ij} has minimal influence on the outcomes of SAM. They have no influence on the detection of candidate policy parameters because control centrality calculations are based on the adjacency matrix only, that is, information related to the connectivity of variables. The impact of τ_{ij} on the path analysis and ADAS methods is minimal because the values only interfere with the total delay of a causal path or a feedback loop and not with the path or loop set detected. So, τ_{ij} only captures information

⁴ To code relationships including a stock variable with $\tau_{ij} = 2$ was explicitly recommended by one of the managing editors of the *System Dynamics Review*—the academic outlet of the System Dynamics Society—in a previous publication project [15].

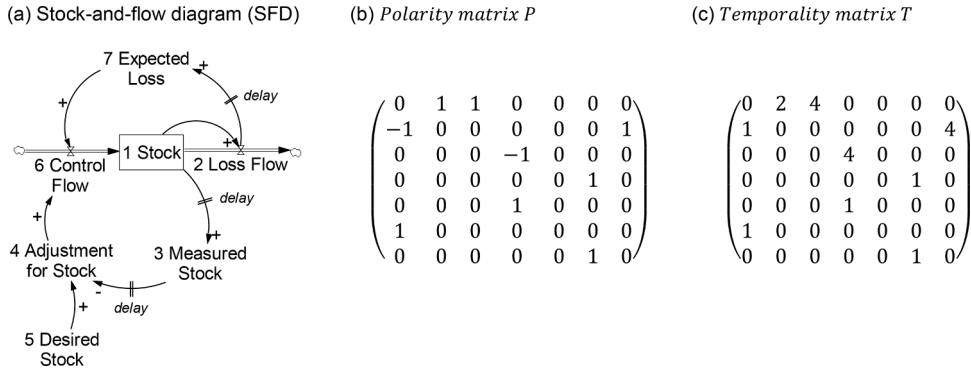


Fig. 6. Stock-and-flow diagram (SFD) with polarity matrix P and temporality matrix T .

relevant to characterise paths or loops but not to identify them.

4.3. Polarity and Relative Delay of Paths

A directed path w_i is a series of disjoint (distinct) connected vertices $v_i \rightarrow v_{i+1} \rightarrow \dots \rightarrow v_n$ where $v_i \neq v_n$. The polarity of a path w_i is calculated by multiplying the path's edge polarities p_{ij} [14, 15]:

$$\text{polarity of path } w_i = \prod p_{i,j} \quad (4)$$

Analogous to the polarity of single edges, a positive path polarity implies that if the cause variable v_i , i.e., the first variable on a path w_i , increases, the effect variable v_n , i.e., the last variable on a path, increases *above what it could otherwise have been* or vice versa, so the change is in the same direction. A negative path polarity implies that if the cause variable v_i at the head of a path w_i increases, the effect variable v_n at the tail of a path decreases *below what it would otherwise have been* or vice versa, so the change is in the opposite direction [28].

The relative delay of a path w_i is calculated by adding the path's edge delays τ_{ij} [15]:

$$\text{relative delay of path } w_i = \sum \tau_{ij} \quad (5)$$

The relative delay of a path w_i measures the relative time elapsing between a change in the cause variable v_i at the head of a path and a change in the effect variable v_n at the tail of a path.

5. Structural Analysis Methods (SAM)

5.1. Detecting candidate policy parameters

The detection of candidate policy parameters in SD models is based on network controllability, a powerful subfield within network science [24]. A system is said to be controllable if it can be directed from an arbitrary initial state to any desired final state in finite time [40]. Controllability can be demonstrated easily with stick balancing, that is, by balancing a stick on one's palm: we know from our experience that stick balancing is possible, inferring that the system must be controllable [41]. In general, controllability is a requirement for control, thus understanding the topology (structure) of the underlying network that determines a system's controllability, provides many insights into the control principles of complex systems [42]. The foundation of network controllability is the linear time-invariant control system:

$$\dot{x}(t) = Ax(t) + Bu(t) \quad (6)$$

where $x(t)$ is a column vector representing the state of the N vertices at time t , A is the adjacency matrix of the underlying network, $B = (b_{i,m})$ is the input matrix identifying the vertices that are directly controlled for, and $u(t)$ is an input vector. The linear time-invariant control system in Equation (6) is controllable if and only if the controllability matrix C has full rank, i.e.:

$$\text{rank}(C) = N \quad (7)$$

where $C = [B, AB, A^2B, \dots, A^{N-1}B]$ [40]. If Equation (7) is satisfied, then there exists an appropriate input vector $u(t)$ to steer the system from any initial state $x(0)$ to any arbitrary final state $x(t)$, implying that the system is controllable [24]. From an inspection of the controllability matrix C it becomes evident that the network structure, captured in A , has a key role in controllability. In large networks, however, determining C is computationally demanding and often not all edges (and edge weights), i.e., the elements in A , are precisely known. To circumvent the latter problem, Liu, Slotine and Barabasi [24] apply the concept of structural controllability [43] where A and B are considered structured matrices, i.e., their elements are either fixed zeros or independent free parameters. The system

in Equation (6) is structurally controllable if we can set the nonzero elements in A and B such that the resulting system satisfies Equation (7) [24].

Liu, Slotine and Barabasi [25] introduced the measure of control centrality to quantify the ability of a single vertex in controlling an entire network. The control centrality of a vertex v_i captures the dimension of the controllable subspace or the size of the controllable subsystem when we control v_i only. When this is the case, B reduces to a vector with a single non-zero entry (corresponding to the node v_i). The resulting controllability matrix C^i is computed as above (see Fig. 7 for an example). The control centrality of v_i is calculated as the rank of the matrix C^i . As A and B are considered structured matrices, instead of the rank of C^i , which is difficult to determine, we can compute the generic rank, $rank_g$. Consequently, the control centrality $C_c(i)$ of a vertex v_i is defined as:

$$C_c(i) = rank_g(C^i) \quad (8)$$

In case $rank_g(C^i) = N$, then vertex v_i alone can control the whole network. Any value of $rank_g(C^i)$ smaller than N describes the dimension of the subspace v_i can control. Notably, if $rank_g(C^i) = 1$, then v_i can only control itself. Equation (8) can be normalized resulting in:

$$c_c(i) = \frac{C_c(i)}{N} \quad (9)$$

To calculate $c_c(i)$ for variables in SD models, the only prerequisite is encoding the model in the form of its adjacency matrix A .⁵ We use the control centrality $c_c(i)$ of a vertex v_i (variable) as a proxy for its influence on model behaviour; that is, the greater the value of $c_c(i)$ is, the more influence a variable has on model behaviour. Schoenenberger and Tanase [12] calculated $c_c(i)$ for all variables in the World2 model [19], and observed that individually changing the highest scoring variables by 10% has a significant impact on the trajectory of all stock variables, while an individual change of 10% to the lowest scoring variables has no impact on the stock variables at all. Thus, evidence suggests $c_c(i)$ to be a reasonable indicator for the detection of candidate policy parameters in SD models.

5.2. Detecting intended and unintended consequences of candidate policy parameters

To detect intended and unintended consequences of policy parameters or in general of all policy variables, path analysis methods can be used. They serve as a tool for causal inference in qualitative SD. Causal inference deals with determining the impact of a given cause variable on a given effect variable [44]. For example, this might be the impact of changing a policy variable (cause variable), typically a parameter, on an outcome variable (effect variable) in a causal loop diagram or SFD [28]. For this purpose, a path analysis is performed, where all paths from a given cause variable to a given effect variable are analysed in terms of their relative impact on the effect variable and the relative magnitude of their delay. The polarity of paths is used to determine a path's relative impact on an effect variable and the relative delay of paths to distinguish between short-term and long-term effects of cause variables (e.g., policy variables) [28].

Table 1 shows two examples of causal paths with their corresponding relative impact (polarity of path) and relative delay. The first path implies that the *Desired Stock*, that is, the cause or policy variable, has a positive relative impact on the *Stock*, that is, the effect or outcome variable. Consequently, if the *Desired Stock* grows, this increases the *Stock* above what it would otherwise have been or vice versa. The relative delay of this path is three because it includes three non-delayed causal relationships only. In contrast, the second path means that if the *Measured Stock* rises, the *Stock* will fall below what it would otherwise have been or vice versa. So, the *Measured Stock* has a negative relative impact on the *Stock*. The relative delay of this path, i.e., the sum of its individual edge delays, is six because it includes one significantly delayed relationship and two non-delayed ones. Both the detection of these paths and the measurement of their relative impact and relative delay values are entirely automated based on an algorithmic approach requiring the *polarity matrix* P and the *temporality matrix* T as input. The algorithm takes the initial variable (cause variable) and searches for all possible paths toward the target variable (effect variable). Each path is unique, and a variable can be crossed only once per path. In complex SD models where hundreds or even thousands of unique paths potentially lie between cause and effect variables, paths can be displayed as frequency distributions [15, 29].

The path analysis method can be used as a preliminary instrument to assess different candidate policy variables because it helps decision-makers to recognize both intended and unintended consequences of triggering these variables. An intended consequence of a policy variable means that the outcome variable behaves (changes) as expected from the viewpoint of the decision-maker (intervenor); whereas an unintended consequence of a policy variable implies that the outcome variable behaves (changes) in an unforeseen and probably counterproductive way. Clearly, any feedback loops in which the outcome variable is embedded will modify the influence of a change in the policy variable. Any balancing loops will reduce the impact of changes of the policy variable on the outcome variable, while reinforcing loops will increase the impact. Such feedback loops will, however, not alter the direction of a change, due to the impact of a policy variable. The path analysis method can therefore correctly detect intended and unintended consequences of policy variables [28].

⁵ Controllability analysis can be carried out using the C++ code package: "CalControlCentrality" 25. Creator permission granted; the C++ code is available under <https://scholar.harvard.edu/yyl/code>.



Fig. 7. Derivation of the controllability matrix used to calculate the control centrality. The left-hand side of the figure represents a simplified dynamical system with four nodes ($v_1 - v_4$). To calculate the control centrality of a node (v_1), we set v_1 as the single node to be controlled. That is, we specify B as a vector with a single non-zero element (b_1), corresponding to v_1 . The controllability matrix is calculated as $C^1 = [B, AB, A^2B, A^3B]$. The control centrality $C_c(1)$ is then calculated as the generic rank of C^1 . Example inspired from Liu, Slotine and Barabasi [24].

Table 1
Relative impact (polarity of path) and relative delay of two different causal paths.

Causal path	Relative impact	Relative delay
Desired Stock \rightarrow_+ Adjustment for Stock \rightarrow_+ Control Flow \rightarrow_+ Stock	+	3
Measured Stock $\xrightarrow{\text{Delay}}_-$ Adjustment for Stock \rightarrow_+ Control Flow \rightarrow_+ Stock	-	6

5.3. Detecting candidate structural changes

ADAS is a tool for the detection of the four generic problem archetypal structures as specified by Wolstenholme [31]: the underachievement, relative achievement, relative control and out-of-control archetypal structures. All four structures are made up of an intended consequence (*ic*) feedback loop and an unintended consequence (*uc*) feedback loop. For example, the relative control archetype is composed of two balancing feedback loops: (i) the first capturing the intended consequence loop that controls a certain problem, and (ii) the second representing the unintended consequence loop which counteracts the behaviour of the first, resulting in relative problem control only (see Table 2).

As the archetypal structures often cause problematic system behaviour, their detection can support system diagnosis and improvement. Thus, ADAS is a tool for answering the following question: *Which archetypal structures potentially cause dysfunctional behaviour of a variable of interest?* Systematic identification of all four generic problem archetypal structures significantly simplifies the design of effective policies, because a solution archetypal structure exists for each problem archetypal structure [31, 15]. Solution archetypes suggest structural changes (i.e., addition of variables, links, loops etc.) that need to be implemented in order to override the negative consequence for model (system) behaviour caused by archetypal structures.

Due to the structural equivalence of three generic problem archetypal structures, ADAS relies on only two different algorithms to discover the four archetypal structures [15], requiring polarity matrix P and temporality matrix T as input. In addition to these two matrices, ADAS requires the specification of an outcome variable (variable of interest) that a modeller intends to control or influence. The outcome variable marks the anchor point to initialize the search for archetypal structures. Consequently, the selection of the outcome variable is key to a meaningful final set of detected archetypal structures.

6. Validation of SAM

In the following, we apply SAM to two models of increasing complexity—the third-order Phosphorus Loops in Soil and Sediment (PLUM) model [18, 45] and the fifth-order World2 model [19]—in order to test and validate the approach. We selected these two test models because they are well known in the SD field and both their structure and behaviour have been extensively analysed, facilitating the validation of results generated by SAM. Also, these models are openly accessible and well documented, i.e., provide full equation listing, enabling model exploration by people not involved in the modelling process.

6.1. Phosphorus Loops in Soil and Sediment (PLUM) model

The Phosphorus Loops in Soil and Sediment (PLUM) model is a simple model of a shallow lake that was first published by Carpenter

Table 2
 Wolstenholme's [31] generic two-loop system archetypes.

		Unintended consequence loop	
		Opposition (B)	Competition (R)
Intended consequence loop	Control (B)	B/B Relative control archetype	B/R Out-of-control archetype
	Growth (R)	R/B Underachievement archetype	R/R Relative achievement archetype

B: balancing feedback loop; R: reinforcing feedback loop.

[45] in the *Proceedings of the National Academy of Sciences of the United States of America* (PNAS). The model was developed to study lake eutrophication, i.e., the process of accumulating excess phosphorus in lakes leading to a rapid growth of algae. The formation of algal blooms causes plant death by blocking sunlight, fish death by generating anoxic conditions, and phosphorus recycling from the lake sediment, which rises the already high levels of phosphorus in the system [46]. Lake eutrophication is not only detrimental for the lake biota and biodiversity, it has also a negative impact on the capacity of lakes to provide ecosystem services. Furthermore, cultural services such as recreation and tourism are negatively affected as well [47].

The PLUM model recreates the key dynamics of lake eutrophication which is a sudden regime shift (critical transition) from a nutrient poor, high biodiversity, clear water state to a nutrient rich, low biodiversity, eutrophic state [18]. This regime shift is not irreversible but returning to the clear water state is challenging due to strong hysteresis loops. i.e., loops in which lags between causes and effects occur. Carpenter [45] conceptualized and parametrized the model according to the Lake Mendota in Wisconsin, USA, using empirical data for soil, lake, and sediment phosphorus levels. Abram and Dyke [18] turned Carpenter's [45] ordinary differential equations model into an SD model which we use as a test model for the application and validation of SAM (see Fig. 8).

We use a model-based hypothesis as a reference point for evaluating and contextualising the results generated by SAM in the next sub-sections. Based on Abram and Dyke's [18] analysis of the PLUM model, the following hypothesis can be deduced:

In the PLUM model, lake eutrophication is driven by a shift in loop dominance⁶, i.e., from a system that is controlled by a balancing feedback loop to one that is dominated by a reinforcing feedback loop.

The balancing feedback loop stabilizes (controls) the level of phosphorus in the water by increasing the phosphorus runoff from the lake (*Poutflow*) when the level of phosphorus in the water starts to rise. This balancing loop is illustrated by variables $9 \rightarrow 10 \rightarrow 9$ in Fig. 8. The reinforcing feedback loop leads to an escalation of phosphorus in the water by the so-called phosphorus recycling mechanism. The phosphorus recycling mechanism accelerates phosphorus accumulation in the water because the higher the concentration of phosphorus in the water the more phosphorus is released from the lake sediment. This reinforcing loop is illustrated by variables $9 \rightarrow 13 \rightarrow 8 \rightarrow 9$ in Fig. 8. In summary, if the inflow of phosphorus into the water reaches a certain tipping point, the reinforcing feedback loop kicks in and starts to override the stabilizing (controlling) impact of the balancing feedback loop—the level of phosphorus in the water gets out-of-control [48].

6.1.1. Detecting candidate policy parameters

For network controllability to be applied, we have encoded the PLUM model in the form of its standard adjacency matrix A . To that end, the model is pre-processed in the following way: (i) we translate the model into a directed digraph; (ii) we number all the variables (nodes) in the network, and (iii) we compile the adjacency matrix with $e_{ij} = 1$ if an edge from i to j exists, and $e_{ij} = 0$ otherwise. The adjacency matrix is then used as an input to calculate the control centrality $c_c(i)$ for every node in the model (see Table 3).

Based on the score, we can test if control centrality $c_c(i)$ is a reliable measure for the influence a variable (node) has on the behaviour of the PLUM model by comparing high scoring variables with low scoring variables. For this comparison, we only focus on a specific subset of all variables: the parameters. From an SD perspective, the modeller can only exert control over the parameters and consequently it is this variable type that is mostly used in policy design.

As can be seen in Table 3, a classification of parameters in the PLUM model based on $c_c(i)$ results in just two parameter groups: four parameters with a $c_c(i)$ value of 0.409091 (highest score) and eight parameters with a $c_c(i)$ value of 0.318182. We now assess if the four parameters from the first group have on average a greater impact on the key stock variable, *Phosphorus Water*, than the eight parameters from the second group. We approximate "impact" by the extent the trajectory of the stock variable *Phosphorus Water* is shifted due to a change in a parameter value. More specifically, we perform the following test:

- 1 We randomly pick four nodes from the lower scoring group (sample size = 8); and compare them against the four higher scoring nodes.
- 2 In the sense of policy experiments, we separately increased the value of all eight nodes by 10% and ran eight different simulations, that is, in every simulation run only one variable is changed.
- 3 We evaluate the impact these changes have on the key stock variable, *Phosphorus Water*, by visual inspection only. As a reference curve, a base run according to Abram and Dyke's [18] model parameterization is carried out.

Panels (a) and (b) in Fig. 9 show the impact of an individual 10% increase of both the four top scoring nodes, i.e., the ones with the

⁶ Loop dominance describes the characteristic of feedback systems in which a loop is strong enough to determine the behavior mode of a part of the system [48].

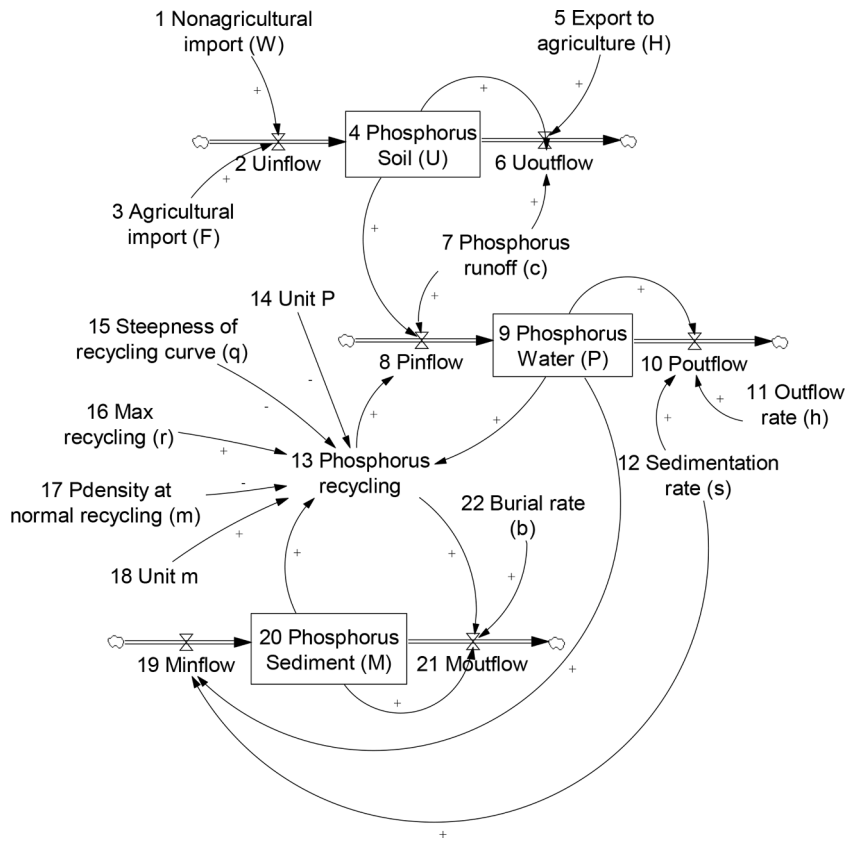


Fig. 8. Phosphorus Loops in Soil and Sediment (PLUM) model.

Table 3
Control centrality $c_c(i)$ of all nodes in the PLUM model (in descending order).

Node name	Type	$c_c(i)$
Export to agriculture (H)	Parameter	0.409091
Phosphorus runoff (c)	Parameter	0.409091
Nonagricultural import (W)	Parameter	0.409091
Agricultural import (F)	Parameter	0.409091
Uoutflow	Flow variable	0.363636
Uinflow	Flow variable	0.363636
Phosphorus Soil (U)	Stock variable	0.318182
Outflow rate (h)	Parameter	0.318182
Sedimentation rate (s)	Parameter	0.318182
Unit P	Parameter	0.318182
Steepness of recycling curve (q)	Parameter	0.318182
Max recycling (r)	Parameter	0.318182
Pdensity at normal recycling (m)	Parameter	0.318182
Unit m	Parameter	0.318182
Burial rate (b)	Parameter	0.318182
Minflow	Flow variable	0.272727
Moutflow	Flow variable	0.272727
Phosphorus recycling	Auxiliary variable	0.272727
Poutflow	Flow variable	0.272727
Pinflow	Flow variable	0.272727
Phosphorus Sediment (M)	Stock variable	0.227273
Phosphorus Water (P)	Stock variable	0.227273

highest $c_c(i)$, and the four lower scoring nodes on the key stock variable *Phosphorus Water*. It is noteworthy that all policy experiments change the trajectory of the reference curve for *Phosphorus Water*. A change in the four lower scoring nodes causes the reference curve to shift in time. An increase of *agricultural import* and *export to agriculture*, two of the top scoring nodes, however, causes the reference curve to significantly change its shape and thus the model's mode of behaviour. For example, this means that if the import of

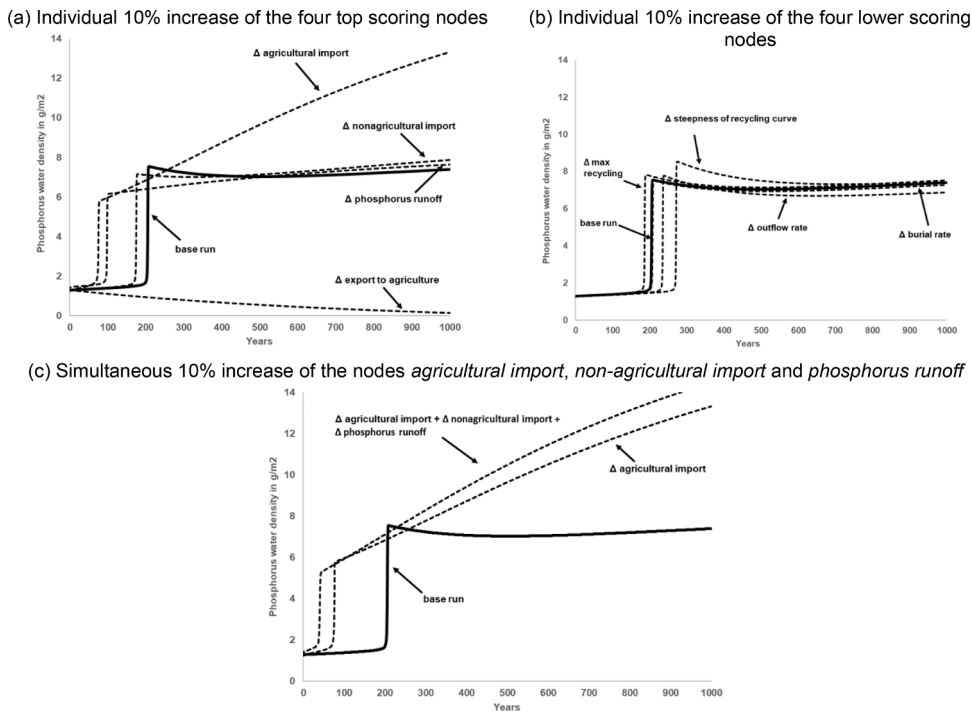


Fig. 9. Impact of individually and collectively changing higher and lower scoring parameter nodes on Phosphorus Water.

phosphorus from agriculture into the watershed of the lake is increased by 10%, the phase transition will still be observable but there is no saturation of phosphorus thereafter—a model behaviour divergent from the reference case. This simple experiment offers evidence supporting the hypothesis that high scoring nodes have more impact on the behaviour of the PLUM model than low scoring nodes.

Besides analysing the impact of changing one parameter value at the time, we might want to examine the effect on the level of phosphorus in the water if several parameters are varied simultaneously. From an inspection of panel (a) in Fig. 9, it becomes evident that among the four top scoring nodes three influence the trajectory of *Phosphorus Water* in a similar way, i.e., *agricultural import*, *non-agricultural import*, and *phosphorus runoff* from the soil into the water. Therefore, it seems reasonable to study the combined impact of changing these three variables together. Panel (c) in Fig. 9 shows the result of such an experimental design. The combined impact of a 10% increase of the three parameter nodes on *Phosphorus Water* is not much larger than the effect of *agricultural import* alone. An extension of SAM to account for the combined effect in the identification of model levers could provide an important avenue for future research.

6.1.2. Detecting intended and unintended consequences of candidate policy parameters

We assess the intended and unintended consequences of triggering the previously identified model levers, that is, *agricultural import* and *export to agriculture* on the key stock variable *Phosphorus Water*. Translated into the context of the PLUM model, this might relate to the following policies: (i) introduction of a new phosphorous poor fertilizer to reduce the import of agricultural phosphorous into the water – a policy reflected in variable 3, *agricultural import*, in Fig. 8; and (ii) a policy of cultivating high phosphorous absorbing plants to increase export of phosphorous from the watershed into farm products – a policy reflected in variable 5, *export to agriculture*, in Fig. 8.

Table 4

Path analysis revealing potential intended and unintended consequences of policies on the level of phosphorous in the lake water (variable 9).

Intended consequence(s) of policy				Unintended consequence(s) of policy			
<i>Introduction of phosphorous poor fertilizers (variable 3)</i>							
No	P	D	Path ^a	No	P	D	Path
1	+	5	3 → 2 → 4 → 8 → 9	1	-	8	3 → 2 → 4 → 8 → 9 → 10 → 9
<i>Cultivating high phosphorous absorbing plants (variable 5)</i>							
1	-	5	5 → 6 → 4 → 8 → 9	1	+	8	5 → 6 → 4 → 8 → 9 → 10 → 9

No: number; P: polarity of path; D: relative delay of path.
 aNumbers correspond to variable numbers in Fig. 8.

As discussed previously, we examine the impact of these policies by performing a path analysis that reveals both their intended and unintended (side) effects on the outcome variable *Phosphorous Water* (variable 9) in the PLUM model. To that end, we compiled the polarity matrix P and the temporality matrix T according to the information in Fig. 8. We assumed that there is not any significantly delayed relationship in the model. The results of the path analysis are summarized in Table 4.

Table 4 shows both policies entail one causal path representing an intended policy consequence and one causal path representing an unintended policy consequence. Consider for example the causal path describing the intended consequence of introducing phosphorous poor fertilizers in agriculture: $3 \rightarrow 2 \rightarrow 4 \rightarrow 8 \rightarrow 9$. This path contains five variables and connects *agricultural import* (variable 3) with *Phosphorous Water* (variable 9). Furthermore, this path has a positive polarity, that is, the lower the import of phosphorous from agriculture into the water – or the more phosphorous poor fertilizers are used, the smaller the amount of phosphorus in the lake water. Finally, the path has a relative delay of five meaning it is a rather short-term effect when compared to the unintended consequence path that exhibits a relative delay of eight. The path reads as follows: introducing phosphorous poor fertilizers (variable 3) reduces the inflow of phosphorous into the soil (variable 2). Consequently, the stock of phosphorous in the soil falls (variable 4) causing the inflow of phosphorous from the soil into the lake water (variable 8) to drop. This over time decreases the level of phosphorous in the water (variable 9).

However, a policy of introducing new phosphorous poor fertilizers also has unintended effects. Consider the causal path depicted in the second column of Table 4: $3 \rightarrow 2 \rightarrow 4 \rightarrow 8 \rightarrow 9 \rightarrow 10 \rightarrow 9$. This path is an extension of the intended consequence path by adding variable 10, the outflow of phosphorous from the water. This means that by introducing phosphorous poor fertilizers (variable 3) the stock of phosphorous in the water drops (variable 9) causing the outflow of phosphorous from the water (variable 10) to lower. As the outflow of phosphorous from the water (variable 10) decreases the stock of phosphorous in the water (variable 9) rises again. In a similar way, the intended and unintended consequences of establishing a new policy that fosters the cultivation of high phosphorous absorbing plants can be interpreted.

6.1.3. Detecting candidate structural changes

As mentioned before, the application of the ADAS method requires the definition of an outcome variable. For the analysis of the PLUM model, we set the key stock variable, *Phosphorous Water* (variable 9), as the outcome variable. More specifically, we want to evaluate if the undesired growth of phosphorous in the lake water, eventually leading to the phase transition from the clear to the turbid water state, is caused by an archetypal structure. Table 5 lists all archetypal structures found in the PLUM model, with *Phosphorous Water* as the outcome variable.

In total, the ADAS method identifies two underachievement archetypal structures, which is not surprising given the low feedback complexity of the PLUM model (six feedback loops). Considering the meaning of an underachievement archetype, that is, intended growth of an outcome variable which is eventually limited by an opposing systemic reaction, it becomes clear that this archetypal structure cannot provide a sensible narrative for the key behaviour observed in the PLUM model. In this case, the intended action would be to control the increase of phosphorous in the lake and not to support it. According to Table 2, if the intended consequence is controlling an outcome variable, only relative control and out-of-control archetypal structures can provide explanations for counterintuitive system behaviour. Since ADAS has not identified any such archetypal structures, the results of this analysis offer little explanatory insights. This is different in the analysis of the World2 model where results produced by ADAS are meaningful in understanding and improving system behaviour.

6.1.4. Evaluation of results

The application of SAM to the PLUM model produced interesting insights. Among the parameter nodes with a relatively higher control centrality two could change the mode of behaviour of the key outcome variable *Phosphorous Water: agricultural import and export to agriculture*. Both model levers were tested for intended and unintended consequences if they were used as policies to lower the level of phosphorus in the water. The key dynamics mentioned in the hypothesis up-front fits well to an out-of-control archetypal structure. However, no such structure was detected by the ADAS method. The reason is that ADAS only identifies structures that precisely conform to Wolstenholme's [31] specification of the four archetypal structures. He specified the out-of-control archetypal structure in such a way that the intended balancing feedback loop and the unintended reinforcing feedback loop must partly overlap, implying that the two loops are not just "held together" by the outcome variable. Therefore, the balancing and reinforcing loops described in the hypothesis up-front do not meet the structural criteria to be classified as an out-of-control archetypal structure in the current version of the ADAS method.

6.2. World2 model

The creation of the World2 model [19] was motivated by the Club of Rome project on the "predicament of mankind" with the primary objective of fostering understanding about the transition from world growth to world equilibrium [49]. The World2 model interconnects concepts from demography, economics, agriculture, and technology and is built around five stock variables: (i) population, (ii) capital investments, (iii) natural resources, (iv) fraction of capital devoted to agriculture, and (v) pollution.⁷

Similarly to the PLUM model, we use in this example a hypothesis as a reference point for assessing and contextualising the results

⁷ An implementation of the World2 model in the modelling software Vensim® can be found in Tom Fiddaman's model library "MetaSD": <https://metasd.com/model-library/> [last accessed 26th Jan 2021]

Table 5
Algorithmically detected archetypal structures.

Intended consequence loop				Unintended consequence loop		
<i>Underachievement archetypes</i>						
No	P	D	Path	P	D	Path
1	+	4	9 → 13 → 8 → 9	-	3	9 → 10 → 9
2	+	7	9 → 19 → 20 → 13 → 8 → 9	-	3	9 → 10 → 9

No: number; P: polarity of loop; D: relative delay of loop.

aNumbers correspond to variable numbers in Fig. 8.

produced by SAM in the next sub-sections. Based on Jantsch's [50] review of the World2 model, the following hypothesis can be made:

To prevent world population from “overshooting and collapsing” investments must be made such that pressures (limitations) on mankind are relaxed.

A key message of World2 was that if the status quo (at the time when the book was published) in terms of population growth and rate of capital investment (industrialisation) continues, natural resources on the planet will be gradually and irreversibly depleted ultimately halting the world system. Here, the natural resources act as a so-called pressure or limitation to the world's growth. In World2, besides the depletion of natural resources other pressures such as pollution, crowding, and food shortage are considered. So, in essence, mankind has two options on how to deal with the limits to growth: (i) push the limits further away, or (ii) keep human activity within the limits. The above hypothesis relates to the former and suggests capital investments should be made to mitigate the impact of pollution, depletion of natural resources, crowding, and food shortage [50].

6.2.1. Detecting candidate policy parameters

Table 6 shows the control centrality $c_c(i)$ of all parameters in the World2 model. Overall, $c_c(i)$ categorizes the World2 parameters into five groups. Analogous to the PLUM model, we compare high scoring parameters with low scoring parameters in terms of their impact on the trajectory of the key stock variable *Population*. We slightly adapt the experimental procedure applied in the PLUM model:

- 1 We randomly pick five nodes from the two groups with the smallest $c_c(i)$ (sample size = 9); and compare them against the five highest scoring nodes.
- 2 In the sense of policy experiments, we separately increased the value of all ten nodes by 10% and ran ten different simulations, that is, in every simulation run only one variable is changed.
- 3 We evaluate the impact these changes have on the key stock variable, *Population*, by visual inspection only. As a reference curve, a base run according to Forrester's (1971) model parameterization is carried out.

Panels (a) and (b) in Fig. 10 shows the effect of a 10% increase of both the five influential nodes, that is, the ones with the highest $c_c(i)$, and of five ineffective nodes, that is, nodes with low $c_c(i)$, on *Population* – the key stock variable in the World2 model. Most notable, only the parameter *capital agriculture fraction normal* from the sample of nodes with a small $c_c(i)$ has an impact on the trajectory of *Population*. In contrast, raising the five highest scoring nodes by 10% has a substantial impact on the trajectory of *Population*. In particular, increasing the parameter *capital investment rate normal* by 10% not only changes the maximum and final equilibrium point of the trajectory but also its general shape (mode of behaviour), that is, from overshoot and collapse to damped oscillation. Again, this experiment seems to support the assumption that control centrality $c_c(i)$ is a useful indicator for the extent of influence a variable has on model in behaviour in SD.

In addition to analysing the impact of changing one parameter value at a time, similarly to the PLUM model, we might want to study the effect of simultaneously changing several parameters on *Population*. For the sake of simplicity and for illustrative reasons, we chose all five top scoring parameter nodes to be included in a more complex experimental set-up. Panel (c) in Fig. 10 shows the corresponding result. Interestingly, the combined impact of a 10% increase of the five parameter nodes completely neutralises the effect that an individual change in *capital investment rate normal* has on *Population*.

6.2.2. Detecting intended and unintended consequences of candidate policy parameters

For the analysis of the intended and unintended policy consequences, we focus on the parameter *capital investment rate normal*, the parameter that exerted the greatest influence on *Population* in the previous experiment. Interpreted in the context of the World2 model, this might relate to the following policy: introduction of a low interest policy to foster capital investment which in turn increases capital flowing into the construction of cities raising the size of the population on the planet. So, the intended consequence of this policy would be the more we invest the more people live on the planet. Owing to the high number of causal paths connecting *capital investment rate normal* with *Population* – 82 paths in total – we refrain from listing every path in tabular form as in Table 4 and use a frequency distribution instead (see Fig. 11).

Fig. 11 shows two frequency distributions: (i) the distribution of causal paths referring to intended policy consequences illustrated by white bars; and (ii) the distribution of causal paths referring to unintended policy consequences illustrated by black bars. For both

Table 6
Control centrality $c_c(i)$ of all parameters in the World2 model (in descending order).

Node name	Type	$c_c(i)$
Population density normal	Parameter	0.4375
Land area	Parameter	0.4375
Capital investment rate normal	Parameter	0.4375
Capital depreciation rate normal	Parameter	0.4375
Natural resource utilization normal	Parameter	0.4375
Capital initial	Parameter	0.421875
Natural resources initial	Parameter	0.421875
Death rate normal	Parameter	0.40625
Birth rate normal	Parameter	0.40625
Pollution per capita normal	Parameter	0.40625
Pollution standard	Parameter	0.40625
Pollution initial	Parameter	0.390625
Food coefficient	Parameter	0.390625
Food per capita normal	Parameter	0.390625
Effective capital ratio normal	Parameter	0.390625
Capital agriculture fraction normal	Parameter	0.390625
Capital agriculture fraction initial	Parameter	0.390625
Capital agriculture fraction adjustment time	Parameter	0.390625
Pollution initial	Parameter	0.390625
Quality of life normal	Parameter	0.03125

Table 7
Algorithmically detected underachievement archetypes from the SILS.

Intended consequence loop				Unintended consequence loop		
<i>Underachievement archetypes</i>						
No	D	P	Loop	D	P	Loop
1*	3	+	1 → 15 → 1	3	-	1 → 2 → 1
2*	3	+	1 → 15 → 1	5	-	1 → 8 → 7 → 2 → 1
3*	3	+	1 → 15 → 1	7	-	1 → 8 → 23 → 17 → 6 → 2 → 1
4	3	+	1 → 15 → 1	7	-	1 → 31 → 30 → 28 → 5 → 2 → 1
5*	3	+	1 → 15 → 1	8	-	1 → 46 → 45 → 57 → 4 → 2 → 1
6*	3	+	1 → 15 → 1	10	-	1 → 46 → 45 → 57 → 18 → 17 → 6 → 2 → 1
7*	3	+	1 → 15 → 1	11	-	1 → 54 → 52 → 49 → 50 → 30 → 28 → 5 → 2 → 1
8*	7	+	1 → 31 → 30 → 28 → 14 → 15 → 1	3	-	1 → 2 → 1
9*	7	+	1 → 31 → 30 → 28 → 14 → 15 → 1	5	-	1 → 8 → 7 → 2 → 1
10*	7	+	1 → 31 → 30 → 28 → 14 → 15 → 1	7	-	1 → 8 → 23 → 17 → 6 → 2 → 1
11*	7	+	1 → 31 → 30 → 28 → 14 → 15 → 1	8	-	1 → 46 → 45 → 57 → 4 → 2 → 1
12*	7	+	1 → 31 → 30 → 28 → 14 → 15 → 1	10	-	1 → 46 → 45 → 57 → 18 → 17 → 6 → 2 → 1
13*	10	+	1 → 31 → 48 → 46 → 45 → 57 → 13 → 15 → 1	3	-	1 → 2 → 1
14*	10	+	1 → 31 → 48 → 46 → 45 → 57 → 13 → 15 → 1	5	-	1 → 8 → 7 → 2 → 1
15*	10	+	1 → 31 → 48 → 46 → 45 → 57 → 13 → 15 → 1	7	-	1 → 8 → 23 → 17 → 6 → 2 → 1
16*	10	+	1 → 31 → 48 → 46 → 45 → 57 → 13 → 15 → 1	11	-	1 → 54 → 52 → 49 → 50 → 30 → 28 → 5 → 2 → 1
17*	11	+	1 → 25 → 24 → 31 → 32 → 21 → 17 → 12 → 15 → 1	3	-	1 → 2 → 1
18*	11	+	1 → 25 → 24 → 31 → 32 → 21 → 17 → 12 → 15 → 1	5	-	1 → 8 → 7 → 2 → 1
19*	11	+	1 → 25 → 24 → 31 → 32 → 21 → 17 → 12 → 15 → 1	8	-	1 → 46 → 45 → 57 → 4 → 2 → 1
20*	11	+	1 → 25 → 24 → 31 → 32 → 21 → 17 → 12 → 15 → 1	11	-	1 → 54 → 52 → 49 → 50 → 30 → 28 → 5 → 2 → 1

No: number; P: polarity of loop; D: relative delay of loop; *: MSILS

distributions, causal paths are counted according to their relative delay. Additionally, a simplifying assumption is made that the strength of all paths is the same. Consider for example, the frequency distribution of causal paths corresponding to the intended policy consequences (white bars). According to Fig. 11, the model lever *capital investment rate normal* activates 42 causal paths with a positive polarity, i.e., intended policy consequences, among them the fastest one impacting *Population* after a relative delay of 9. However, the impact of this path is compensated by a path with a negative polarity, i.e., an unintended policy consequence, arriving at *Population* after the same relative delay of 9. Thus, we can infer from Fig. 11 at what times (relative delays) an intended impact on *Population* becomes measurable. It is at relative delays of 10, 17, 19, 20, 22, 23, and 34. Overall, because the ratio of causal paths with positive polarity to causal paths with negative polarity is approximately 1, the total effect of increasing the *capital investment rate normal* on

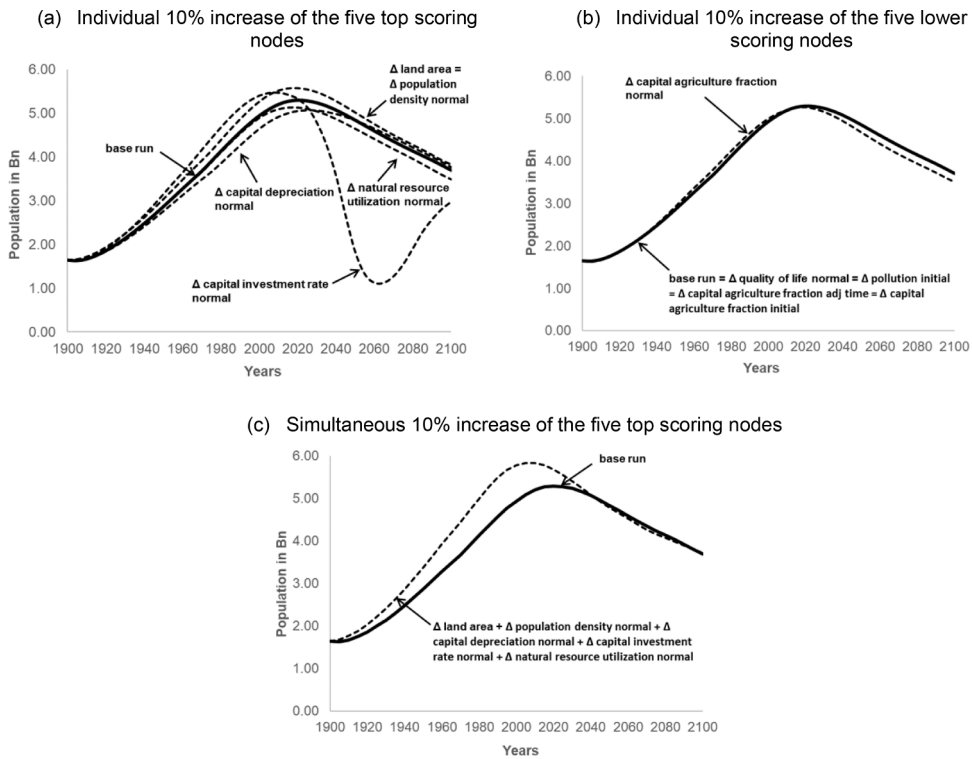


Fig. 10. Impact of individually and collectively changing higher and lower scoring parameter nodes on Population (adapted from [12]).

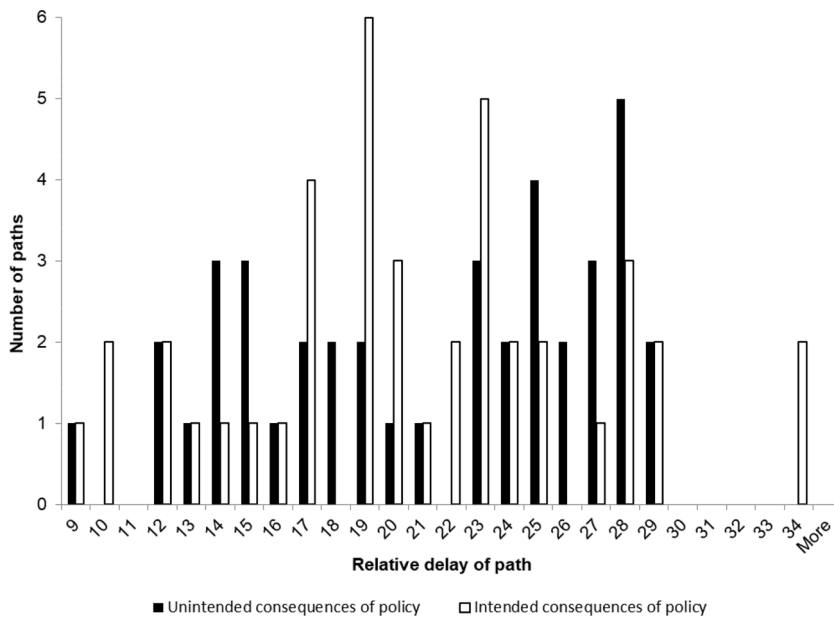


Fig. 11. Frequency distribution of paths referring to intended and unintended policy consequences.

Population according to the path analysis is rather low. This finding corresponds well with population exhibiting oscillatory behaviour when the capital investment rate normal is increased by 10% in the previous sub-section.

6.2.3. Detecting candidate structural changes

As outcome variable for the ADAS analysis, we choose the key stock variable Population. More precisely, we want to assess if the

reference behaviour of population overshoot and collapse (see Fig. 10) is caused by an archetypal structure. Due to the high number of archetypal structures detected in the World2 model, we only focus on the underachievement archetypes; these structures are promising candidates for explaining the reference behaviour of population overshoot and collapse, because they can offer a narrative for both the growth (overshoot) and the decay (collapse) phase. The growth phase can be explained by the reinforcing intended consequence loop that accelerates population growth and the balancing unintended consequence loop that constrains population growth because of approaching resource limits. When both loops operate in the other direction, the decline phase can be described by the reinforcing intended consequence loop that generates exponential decay and the unintended consequence loop that ultimately slows down population decay.

Applying the ADAS method to the World2 model results in a total of 99 underachievement archetypes. Obviously, this high number of identified underachievement archetypes prevents modellers from efficiently designing impactful policies. Therefore, we propose to work with a reduced loop set instead of algorithmically processing all feedback loops in the model [15]. In SD, three different types of reduced loop sets have been described: (i) independent loop set (ILS); (ii) shortest independent loop set (SILS); and (iii) minimal shortest independent loop set (MSILS) [14, 13].

The application of the SILS algorithm to the World2 model results in a 70% reduction in the total number of feedback loops, that is, the 80 feedback loops in the initial model are reduced to 24 feedback loops in the SILS. Using only the SILS as an input for the ADAS analysis diminishes the resulting archetypes set to 20 underachievement archetypal structures (see Table 7). In contrast, the application of the MSILS algorithm to the World2 model generates a 75% reduction in the total number of feedback loops, i.e., the 80 feedback loops in the initial model are reduced to 20 feedback loops in the MSILS. Using these 20 feedback loops for the ADAS analysis results in a final set of 19 underachievement archetypes (see archetypes marked with one asterisk in Table 7).

Among the 20 underachievement archetypes identified from the SILS, four archetypal structures describe situations in which fostering capital investment is the policy approach adopted to support population growth (archetypes no. 17-20 in Table 7). So, in all these four archetypal structures the reinforcing intended consequence loop tells the following story: The higher capital investment is, the more money being spent on agriculture (i.e., the capital ratio agriculture increases), which in turn rises both the food ratio and the number of births, eventually leading to a larger Population. The four archetypal structures differ in that they describe four different unintended consequences when implementing this policy. Fig. 12 summarizes these four underachievement archetypes showing the common intended consequence loop and the four different unintended consequence loops.

To mitigate the unintended consequences of underachievement archetypes, Wolstenholme [31] suggests introducing so-called “solution links”:

The closed-loop solution to an underachievement archetype lies in trying to use some element of the achievement action to minimise the reaction in other parts of the organisation, usually by unblocking the resource constraint. That is to introduce a further reinforcing loop in parallel with the ic reinforcing loop to counter the balancing reaction.

If we apply this idea to the above policy example of heightened capital investment, it implies that capital investments should be used to minimize the impact of the four unintended consequence loops. For example, this could be to invest capital in the reduction of the

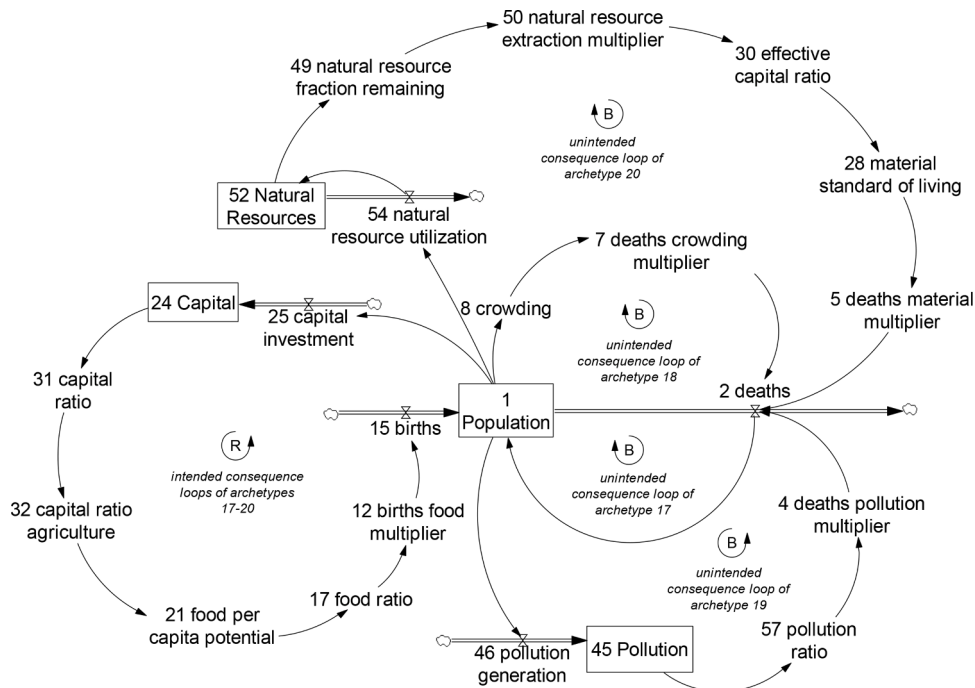


Fig. 12. Underachievement archetypes resulting from a policy of more capital investment (adapted from [15]).

number of *deaths* for the unintended consequence loops 17 and 18, to invest capital in the reduction of the *pollution generation* for the unintended consequence loop 19, and to invest capital in the reduction of the *natural resource utilization* for the unintended consequence loop 20 (see Fig. 12). More specifically, we can implement three solution links with negative polarity from *capital investment* to *deaths*, from *capital investment* to *pollution generation*, and from *capital investment* to *natural resource utilization*. As expected, these structural changes have a strong impact on population growth compared to the initial model, leading to an almost linear growth until the end of the simulation period (see Fig. 13).

6.2.4. Evaluation of results

The application of SAM to the World2 model generated insights consistent with the hypothesis presented up-front. Control centrality calculations revealed five parameters as candidate model levers: *population density normal*, *land area*, *capital investment rate normal*, *capital depreciation rate normal*, and *natural resource utilization normal*. Individually raising the five parameters by 10% shows a notable impact on the trajectory of population. Particularly, increasing the parameter *capital investment rate normal* by 10% changes the mode of behaviour of population from overshoot and collapse to damped oscillation. The path analysis discovered the intended and unintended consequences of implementing a policy of promoting capital investment on population growth. However, this policy alone without any structural changes seems to have a minor total effect on population because the ratio of causal paths with positive polarity (intended consequences) to causal paths with negative polarity (unintended consequences) is approximately one. This observation corresponds well with the fact that the population trajectory starts to oscillate when the *capital investment rate normal* is increased by 10%. Finally, ADAS found four underachievement archetypal structures in the model that cause a policy of promoting capital investment to be an “underachiever”. The four underachievement archetypal structures can be disarmed by introducing structural changes into the model. More specifically, the following three causal links can be implemented to neutralise the archetypal structures:

- 1 A causal link with negative polarity from *capital investment* to *deaths*. This corresponds to a policy of making capital investments to reduce mortality (e.g., improved health care).
- 2 A causal link with negative polarity from *capital investment* to *pollution generation*. This corresponds to a policy of making capital investments to reduce pollution emissions (e.g., new technology)
- 3 A causal link with negative polarity from *capital investment* to *natural resource utilization*. This corresponds to a policy of making capital investments to reduce natural resource utilization (e.g., new technology)

7. Limitations

SAM has a couple of limitations that are worth describing here. A crucial restriction is that SAM only considers structural information of SD models, that is, the number of variables and their interconnections, together with polarities of model relationships and delays. This leads to a substantial simplification of SD models because information on the parameter values and on the mathematical formulation of relationships is lost. We are aware of SD being most interested in system behaviour and not in structure per se. However, one of the key assumptions of SD emphasises that system behaviour arises out of its underlying system structure [13]. Consequently, having a suite of tools at hand that enables a comprehensive analysis of system structure, will eventually lead to an improved understanding of system behaviour. It is important to note, however, that SAM has only been tested on two SD models, albeit of different complexity. So, while results generated with these two models are promising, they are still on a case study level. Therefore, in the future, it is necessary to apply SAM to a multitude of SD models to gain a better understanding on the validity of results.

Referring to the detection of candidate policy parameters, network controllability provides information on which variables (parameters) to focus, in order to attain full model control but it does not specify in which direction and by how much these variables must be changed. Consequently, variables may have to be changed by such an extent that implementation of these changes in the real world is infeasible [12]. Furthermore, when a policy requires controlling $p > 1$ input parameters, selecting the first p nodes with the highest control centrality is not always optimal as there can be an overlap between the nodes controlled by each of the p driver nodes. A straightforward solution is to extend the definition of control centrality to a group of nodes, rather than to a single node [42]. However, this does not directly tackle the problem of identifying the optimal set of driver nodes. Lo Iudice, Garofalo and Sorrentino [51] propose a framework to maximise the diffusion of control signals in a network, while allowing for three types of constraints: i) which nodes can be selected as drivers (*admissible*); ii) which nodes must be controlled (*targets*) and which nodes must not be perturbed in the control process (*untouchables*). This framework allows solving two complementary problems. The first is identifying the set of driver nodes of fixed cardinality, out of the admissible nodes, that maximises the number of target nodes while not perturbing the untouchable nodes. The second is identifying the set of driver nodes of minimum cardinality, out of the admissible nodes, that can control at least the target nodes, while not perturbing the untouchable nodes. Extending the current framework to account for these insights might prove a valuable avenue for further research.

While the path analysis allows to detect the intended and unintended consequences of triggering policy parameters, it does not allow to measure their overall impact and therefore to assess their effectiveness (policy effectiveness). To measure the overall impact of activating policy parameters would still require a fully-fledged simulation model [28].

The ADAS method is not yet capable of identifying “true” system archetypes, only modelling components that fulfil their structural requirements. True system archetypes are more than simple two-loop constellations; they are real-world phenomena with causes and effects separated in time and space [31]. Particularly, spatial differences between the intended and unintended consequence loops delineated by system boundaries are not integrated into the ADAS method. For this reason, modellers’ judgments are still needed for the interpretation of the analysis results, preventing fully automated detection of system archetypes. Additionally, the results

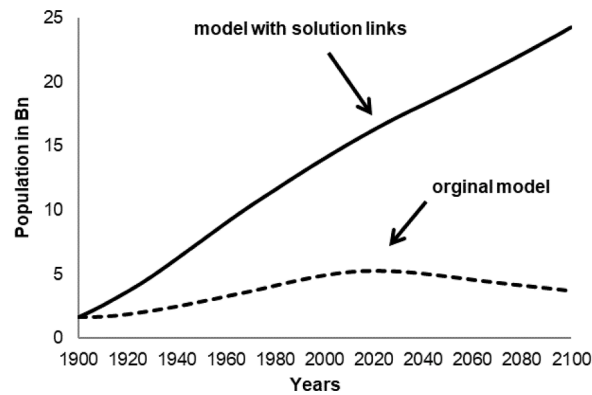


Fig. 13. Population growth in original model and adapted model with solution links (adapted from [15]).

generated by ADAS might be difficult to interpret in complex SD models, because of the current lack of criteria other than structural requirements for the identification of system archetypes. This problem recurs in any complex model because the probability of detecting archetypal structures rises as the number of feedback loops in the model grows. In reality, however, not all of the detected archetypal structures are plausible explanations for counterintuitive system behaviours [15].

8. Discussion and Conclusions

In this paper, we propose a coherent analysis framework, called SAM, in order to facilitate the policy design process. The acronym stands for *structural analysis methods*, subsuming multiple analytical methods under one umbrella term. SAM complements established analytical methods such as pathway participation metric (PPM), model structure analysis (MSA), or eigenvalue elasticity analysis (EEA). Thereby, SAM increases the analytical accessibility of model structure as a foundation for the explanatory analysis of model behaviour, requiring relatively less resources in time and knowledge than traditional, powerful, yet expensive, quantitative analysis techniques. The central assumption of SAM is that the structure of SD models can be accurately described through directed weighted digraphs and underlying matrices, making it accessible for algorithmic exploration by means of graph theory and network science. This requires only little formalization of causal loop or stock and flow diagrams.

SAM in model analytics consists of three steps. Firstly, candidate policy parameters in models are extracted. For that purpose, we suggest using network controllability calculations. Secondly, causal paths are traced to explore model behaviour when triggering policy parameters. Those causal paths may describe intended and unintended consequences of interventions at policy parameters. Thirdly, archetypal combinations of feedback loops are identified and analysed to create an explanatory foundation of model behaviour based on structure. Detecting archetypal structures is a valuable exercise both for system (model) diagnosis and improvement (policy design). It is useful for system diagnosis because archetypal structures can be the cause for problematic dysfunctional system behaviour; and it is helpful for system improvement because for every archetypal structure specific structural changes are recommended in the literature in order to mitigate the negative impact of these structures.

The formal ground on which SAM rests is graph-theoretical notation. Causal maps, causal loop, or stock and flow diagrams can be translated into adjacency matrices. Depending on the information available in the initial models, adjacency matrices serve to describe derived or estimated polarity and temporality of relationships. Polarity and temporality matrices are the required input for the path analysis and the ADAS method. For calculating control centralities, only the standard adjacency matrix is needed which captures the network structure of SD models.

We validate SAM on the basis of two classical, popularly used and cited SD models: the third-order PLUM model by Abram and Dyke [18] (based on Carpenter [45]) and the fifth-order World2 model by Forrester [19]. The PLUM model describes the dynamics of accumulating excess phosphorus in shallow lakes, which in excess leads to a rapid growth of algae, threatening the balance of the whole ecosystem. The model is described by 22 vertices and 29 edges. The World2 model revolves around the dynamics of the transition from global population growth to global equilibrium, interconnecting concepts from demography, economics, agriculture, and technology. The model comprises 64 vertices and 93 edges. Due to their different size and complexity, processing the PLUM and World2 model with SAM yields different analytical insights, allowing to partly describe the behaviour of the models based on their structure.

By analysing the PLUM model with SAM, the results immediately reveal two candidate policy parameters to control the problematic phosphorus concentration in the water: introduction of phosphorous poor fertilizers and cultivation of phosphorous absorbing plants. As highly plausible interventions, simulations show that altering those levers by 10% causes the phosphorus concentration to develop significantly different from the base run. A path analysis uncovers the operating mechanisms behind this change in behaviour, taking effect through the phosphorus concentration in the soil. The final step of SAM does not lead to the identification of archetypal structures that could offer plausible explanations for the key dynamics underlying the process of lake eutrophication. The key dynamics is characterised by a shift in loop dominance, that is, from a system where the level of phosphorus in the lake is stabilised by a balancing feedback loop to a system where the phosphorus concentration suddenly explodes due to an operating reinforcing feedback

loop. In principle, the combination of these two loops would match well with an out-of-control archetypal structure. However, in the current implementation of ADAS, the structural preconditions for a two-loop combination to be classified as archetypal are rather narrow and strictly follow Wolstenholme's [31] definition. It might be beneficial to relax these structural criteria in the future to expand the number of detected archetypal structures.

The analysis of the World2 model with SAM uncovers a relatively large number of intervention possibilities. Based on the control centrality, interventions in capital investment towards agricultural productivity are promising because they not only change the maximum and final equilibrium point of the population trajectory, but also its general shape. However, SAM shows that a large amount of feedback loops balances population growth and decay, even with increased agricultural productivity. Those limits are potentially explained by the 99 underachievement archetypal structures in the model. If an ultimate goal was to remove those limits to population growth, we find with SAM that capital investments into health (against death), resource utilization, and against pollution would be the most immediately effective policy solutions.

Our study shows that SAM offers a productive way of combining existing analytical tools. Based on the insights gained in the PLUM and World2 model, SAM seems to be most valuable for the analysis of more complex models. It simplifies their structural complexity, offering revelatory and plausible insights, gained in a resourceful way. The usefulness of SAM is limited as regards relatively smaller models. Such smaller models are not only easier to be investigated manually, but they also have fewer feedback loops serving as structural explanations for model behaviour. Consequently, the algorithmic composition of the structural analysis methods combined in SAM is best used when qualitative, visual mapping and diagramming methods reach their limits due to the growing incomprehensibility with rising model size.

Acknowledgment

The work in this article was inspired by the ideas of Professor Andrea Schenker-Wicki, currently President of the University of Basel in Switzerland. She has an enormous passion for systems thinking and is particularly interested in combining concepts from system dynamics and cybernetic analysis. Andrea Schenker-Wicki's academic work in the systems area significantly shaped and refined the methodological viewpoint advocated in this article. Furthermore, we would like to acknowledge that the work of Radu Tanase was supported by the University Research Priority Program "Social Networks" of the University of Zurich.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.simpat.2021.102333](https://doi.org/10.1016/j.simpat.2021.102333).

References

- [1] H Rutter, N Savona, K Glonti, J Bibby, S Cummins, DT Finegood, The need for a complex systems model of evidence for public health, *The Lancet* 390 (2017) 2602–2604.
- [2] M Egan, E McGill, T Penney, PS Meier, N Savona, F de Vocht, Complex systems for evaluation of public health interventions: a critical review, *The Lancet* 392 (2018) S31.
- [3] J-A Atkinson, A Page, A Prodan, G McDonnell, N Osgood, Systems modelling tools to support policy and planning, *The Lancet* 391 (2018) 1158–1159.
- [4] WR Stahel, The circular economy, *Nature* 531 (2016) 435–438.
- [5] T Rebs, M Brandenburg, S Seuring, System dynamics modeling for sustainable supply chain management: A literature review and systems thinking approach, *Journal of Cleaner Production* 208 (2019) 1265–1280.
- [6] SR Golroudbary, SM Zahraee, System dynamics model for optimizing the recycling and collection of waste material in a closed-loop supply chain, *Simulation Modelling Practice and Theory* 53 (2015) 88–102.
- [7] J. Sterman, *Business dynamics. Systems thinking and modeling for a complex world*. Maidenhead, Berkshire, McGraw Hill Education, 2000.
- [8] JW Forrester, *Industrial dynamics*, The M.I.T. Press, Cambridge Mass, 1969.
- [9] EG Pagoni, G Patroklos, A system dynamics model for the assessment of national public-private partnership programmes' sustainable performance, *Simulation Modelling Practice and Theory* 97 (2019), 101949.
- [10] JW Forrester, System dynamics, systems thinking, and soft OR, *SystDynRev* 10 (1994) 245–256.
- [11] R Oliva, Structural dominance analysis of large and stochastic models, *SystDynRev* 32 (2016) 26–51.
- [12] L Schoenenberger, R Tanase, Controlling complex policy problems: A multimethodological approach using system dynamics and network controllability, *Journal of Simulation* 12 (2018) 16–170.
- [13] R Oliva, Model structure analysis through graph theory: partition heuristics and feedback structure decomposition, *SystDynRev* 20 (2004) 313–336.
- [14] CE Kampmann, Feedback loop gains and system behavior, 1996, *SystDynRev* 28 (2012) 370–395.
- [15] L Schoenenberger, A Schmid, M Schwaninger, Towards the algorithmic detection of archetypal structures in system dynamics, *SystDynRev* 31 (2015) 66–85.
- [16] C Eden, Analyzing cognitive maps to help structure issues or problems, *European Journal of Operational Research* 159 (2004) 673–686.
- [17] G Montibeller, V Belton, Causal maps and the evaluation of decision options—a review, *Journal of the Operational Research Society* 57 (2006) 779–791.
- [18] JJ Abram, JG Dyke, Structural Loop Analysis of Complex Ecological Systems, *Ecological Economics* 154 (2018) 333–342.
- [19] JW. World Forrester, *Dynamics*. Cambridge, Wright-Allen Press, 1971.
- [20] JW. Forrester, *Industrial Dynamics*, MIT Press, Cambridge, Mass, 1961.
- [21] JW Forrester, *Industrial Dynamics—After the First Decade*, *Management Science* 14 (1968) 398–415.
- [22] GP Richardson, Reflections on the foundations of system dynamics, *SystDynRev* 27 (2011) 219–243.
- [23] GP. Richardson, *Feedback thought in social science and systems theory*, University of Pennsylvania Press, Philadelphia, 1999.
- [24] Y-Y Liu, J-J Slotine, A-L Barabási, Controllability of complex networks, *Nature* 473 (2011) 167–173.
- [25] Y-Y Liu, J-J Slotine, A-L Barabási, Control centrality and hierarchical structure in complex networks, *PLoS ONE* 7 (2012) e44459.
- [26] R Tarjan, Depth-First Search and Linear Graph Algorithms, *SIAMJComput* 1 (1972) 146–160.
- [27] M. Hürlimann, *Dealing with Real-World Complexity*, Gabler, Wiesbaden, 2009.

- [28] LK Schoenenberger, S Bayer, JP Ansah, DB Matchar, RL Mohanavalli, SS Lam, Emergency department crowding in Singapore: Insights from a systems thinking approach, *SAGEOpenMed* 4 (2016), 2050312116671953.
- [29] L Schoenenberger, A Schenker-Wicki, M Beck, Analysing Terrorism from a Systems Thinking Perspective, *Perspectives on Terrorism* 8 (2014) 16–36.
- [30] L Schoenenberger, A Schmid, J Ansah, M Schwaninger, The challenge of model complexity: improving the interpretation of large causal models through variety filters, *SystDynRev* 33 (2017) 112–137.
- [31] EF Wolstenholme, Towards the definition and use of a core set of archetypal structures in system dynamics, *SystDynRev* 19 (2003) 7–26.
- [32] PM. Senge, *The fifth discipline. The art and practice of the learning organization*, Currency Doubleday, New York, N.Y., 2006.
- [33] P Gonçalves, Behavior modes, pathways and overall trajectories: eigenvector and eigenvalue analysis of dynamic systems, *SystDynRev* 25 (2009) 35–62.
- [34] CE Kampmann, R Oliva, Loop eigenvalue elasticity analysis: three case studies, *SystDynRev* 22 (2006) 141–162.
- [35] M Saleh, R Oliva, CE Kampmann, PI Davidsen, A comprehensive analytical approach for policy analysis of system dynamics models, *European Journal of Operational Research* 203 (2010) 673–683.
- [36] M Mojtahedzadeh, Do parallel lines meet? How can pathway participation metrics and eigenvalue analysis produce similar results? *SystDynRev* 24 (2008) 451–478.
- [37] M Mojtahedzadeh, Consistency in explaining model behavior based on its feedback structure, *SystDynRev* 27 (2011) 358–373.
- [38] M Mojtahedzadeh, D Andersen, GP Richardson, UsingDigest to implement the pathway participation method for detecting influential system structure, *SystDynRev* 20 (2004) 1–20.
- [39] TH. Cormen, *Introduction to algorithms*, MIT Press, Cambridge, Mass, 2009.
- [40] RE Kalman, Mathematical Description of Linear Dynamical Systems, *Journal of the Society for Industrial and Applied Mathematics Series A Control* 1 (1963) 152–192.
- [41] DG. Luenberger, *Introduction to dynamic systems. Theory, models, and applications*, New York, 1979 a.o.: Wiley.
- [42] Y-Y Liu, A-L Barabási, Control principles of complex systems, *RevModPhys* 88 (2016) 247.
- [43] C-T Lin, Structural controllability, *IEEETransAutomatContr* 19 (1974) 201–208.
- [44] RM Axelrod, *Structure of decision. The cognitive maps of political elite*, Princeton University Press, Princeton, New Jersey, 1976.
- [45] SR Carpenter, Eutrophication of aquatic ecosystems: bistability and soil phosphorus, *ProcNatlAcadSciUSA* 102 (2005) 10002–10005.
- [46] M Scheffer, SR Carpenter, Catastrophic regime shifts in ecosystems: linking theory to observation, *Trends in Ecology & Evolution* 18 (2003) 648–656.
- [47] WK Dodds, WW Bouska, JL Eitzmann, TJ Pilger, KL Pitts, AJ Riley, Eutrophication of U.S. freshwaters: analysis of potential economic damages, *EnvironSciTechnol* 43 (2009) 12–19.
- [48] DN Ford, A system dynamics glossary, *SystDynRev* 35 (2019) 369–379.
- [49] L Schoenenberger, A. Schenker-Wicki, Can system dynamics learn from social network analysis?, in: *In: Proceedings of the 2014 International System Dynamics Conference. Delft: the Netherlands: System Dynamics Society, 2014.*
- [50] E Jantsch, World dynamics, *Futures* 3 (1971) 162–169.
- [51] F Lo Iudice, F Garofalo, F Sorrentino, Structural permeability of complex networks to control signals, *NatCommun* 6 (2015) 8349.