MARVEL
-
principles of a method for semi-qualitative system behaviour and policy analysis

Erik J.A. van Zijderveld

TNO Defence, Security and Safety
P.O. Box 96864, 2509 JG The Hague, The Netherlands
Tel. +31-70-3740129, Fax +31-70-3740642
erik.vanzijderveld@tno.nl

Abstract

Obtaining insight into the effects of policy interventions is often a difficult matter. A new method to obtain a first insight into those effects is presented in this paper. The basis of the method is a Causal Loop Diagram to which information on causal relations and variables is added. Part of the information is expressed in qualitative terms.

This Method to Analyse Relations between Variables using Enriched Loops (MARVEL) takes proposed interventions as a starting point. Interventions are interpreted as imposed changes on selected model variables representing intervention points. A new feature is that causal relations are no longer passive but active model elements. They propagate the changes through the model in a time-dispersed way. MARVEL determines how this causes (other) variables representing the model's performance to change in the desired direction at selected moments in time.

MARVEL can be used for policy development, policy analysis and policy evaluation problems.

Key words

Causal Loop Diagram; Qualitative Decision Analysis; Policy Analysis; System Dynamics

1 Introduction

The field of System Dynamics offers several modelling and analysis techniques. One is the Causal Loop Diagram (CLD) which is purely qualitative and serves well to gain insight into the problem's structure without requiring much data. Another is the Stock and Flow Diagram (SFD) that generally is quantified and therefore excellent to gain insight into the problem's time behaviour based on numeric formulas and data *(Sterman 2000)*.

For some situations neither a CLD nor an SFD seem ideal. It is recognized that a CLD is imperfect for many situations as it offers limited analysis opportunities *(Coyle 2000)*. Proper use of an SFD is not always feasible as it may require about 10 to 100 times the effort required for a CLD *(Coyle 2001)*. This effort is not acceptable when a decision must be made quickly or when the stakes are not high enough to justify such an investment *(Homer et al. 2001)*.

One of the situations where neither a CLD nor an SFD seem ideal is when a group of stakeholders are exploring policy intervention options for a problem with limited quantitative data available. Practical experience with Group Model Building (GMB) *(Vennix 1996)* sessions on these problems showed to the author that stakeholders may not always be able to produce quantitative data at an acceptable effort but still often can agree on a somewhat holistic view on the "strength" and "speed" of causal relations in a CLD.

For these situations it would be desirable to have a method that can be used in a GMB setting, that handles general input on the "strength" and "speed" of causal relations and that can provide results that generally are relevant for a first policy intervention analysis with several stakeholders. These results often include:

1. insight into and consensus on the problem structure;
2. insight into the relative power of identified feedback loops;
3. first insight into the problem's behaviour over time;
4. first insight into the expected effects of proposed interventions;
5. first insight into the better interventions to propose.

The word 'first' stresses that outcomes do not need to have the accuracy that can be obtained from a quantified SFD[1].

A method meeting these requirements appears to be nonexistent as discussed in Section 2. This paper therefore proposes a new Systems Dynamics related analysis method to perform such a first policy intervention analysis for problems:

1. that can be described in terms of causal relations;
2. that are expected to contain feedback loops;
3. that are expected to show dynamic behaviour;
4. that are about intervention evaluation;
5. for which general ideas on the strength and speed of causal relations exist;
6. for which only a first insight into time behaviour is acceptable;
7. where insight into the relative power of identified feedback loops is desired.

The remainder of this paper is organised as follows. Section 2 provides a short survey of methods found in the literature dealing with similar problems, including their strengths and limitations. Section 3 gives the modelling approach of the new method that was developed to fill the identified gap in currently available methods. That method offers two analysis types of which the principles are described in Section 4. An illustrative

[1] Reasons not to use a quantified SFD may include: limited outcome accuracy is accepted; sheer unavailability of information; too great duration, capacity or other investments required to make the information available, whether in absolute terms or just at the current state of problem formulation.

example is presented in Section 5, followed by a discussion of several practical and methodological issues in Section 6 and concluding remarks in Section 7.

2 Some options for semi-qualitative analysis

The sought for analysis method should offer more (preferably qualitative) analysis opportunities than a CLD and should require less effort than a quantified SFD. It also should be usable in a GMB context *(Vennix 1996)* as typically policy decision analysis involves several stakeholders who ideally should reach consensus on the problem structure. The search for System Dynamics related analysis techniques that are more qualitative than an SFD dates back to at least as early as 1983. In that year Wolstenholme and Coyle describe a first qualitative analysis approach *(Wolstenholme et al. 1983)*. More recent work shows that the quest for additional analysis methods is still going on.

Liddell for instance categorises all feedback loops in a Qualitative Politicised Influence Diagram (QPID) based on type, speed and strength *(Liddell et al. 2004)*. Liddell then pragmatically selects the most promising loops for further analysis to obtain an acceptably timely result. Jac A. M. Vennix' group at the Methodology Department of the University of Nijmegen uses a comparable approach as learned from private discussions with one of Vennix's co-workers. In addition to assigning an overall strength and speed to each feedback loop, also control and goal variables in the CLD are identified. Then the overall impact of each control variable on the entire system is established by mentally combining the effects of the various feedback loops that control variable is part of. For simpler models these methods can deliver good results. For many a real-life problem however several limitations become apparent:

1. mentally combining the effects of various interacting feedback loops often becomes too complex for most humans;
2. parts of the diagram that are common for multiple loops may be appraised and included inconsistently in the speed and strength of each of those loops;
3. interactions between loops themselves are still hard to establish;
4. limited value ranges for variables are not considered[2];
5. often only a subset of the loops is considered for an acceptably timely result.

McLucas proposes a System Dynamics 'Front End' tool *(McLucas 2001)*. It calculates the influence propagation of a pre-selected variable through a CLD, resulting in the final values of the variables. This tool however does not handle speeds or strengths of arrows, nor does it support analysis of interventions on multiple variables in the CLD.

Some recent software tools offer new analysis possibilities. CONSIDEO 2.0 for example allows the user to draw a CLD and add qualitative information to arrows on their sign, speed and impact (see http://www.consideo.de). It also can simulate the time behaviour of the variables. But unfortunately the qualitative information on the arrows' sign and impact is not used in that quantitative simulation. This limited integration of qualitative modelling features and simulation forces to develop a quantitative simulation

[2] A loop may not function if one or more variables in the loop can not change further in the desired direction.

mostly separate from the qualitative model. In the author's view this reduces the added value of that approach. Also the arrows' speed simply represents a fixed delay which is not always realistic and no support for model performance measures is available.

The software tool GAMMA 4.0 allows specifying the relative importance of all input arrows for a variable in a CLD (see http://www.topsim.com/de/vernetztes_denken). The tool positions each CLD variable in a matrix indicating its influence on other variables, and the degree by which it is influenced itself. It however does not include the sign or speed of an arrow, and does not analyse how variables' values develop over time.

Traditional System Dynamics tools like Vensim and Powersim allow constructing a CLD and an SFD. But as stated above in some cases a CLD offers too limited analysis opportunities while an SFD requires too much effort. These traditional tools also do not offer a standardized approach to easily model arrows with their respective speeds and strengths in a GMB setting.

This overview lead to the conclusion that none of the existing approaches, known to the author, would be suitable for the type of policy assessment problem that we were faced with.

3 Modelling approach of MARVEL

The method of Jac A. M. Vennix as learned from private discussions with one of his co-workers and described in Section 2 serves as a basis for the new method as it comes closest to the aimed new method. The idea of control and goal variables is adopted from Vennix's approach as it fits well into the goal of assessing intervention effectiveness. The idea to assign a strength and speed to feedback loops is not adopted as this may lead to inconsistent characterisation of loops that have common sections of the CLD as pointed out earlier. Several new features were added as described in this section.

The resulting new method is called Method to Analyse Relations between Variables using Enriched Loops (MARVEL).

3.1 Starting point

Several observations serve as a guidance for the features to include in the method:

1. A CLD is often a good tool to merge information from various stakeholders about a problem requiring an intervention *(Vennix 1996)*.
2. Practical GMB sessions learn that although quantitative data may be hard to obtain, decision makers can often reach consensus on the speed and strength of causal relations.
3. Decision makers can generally identify potential pressure points and performance indicators in a CLD.
4. It often takes some time before the influence of a change over a causal relation is fully apparent, but before that moment generally some initial effects can be found.

3.2 General approach

MARVEL assumes a model structure consisting of causal relations between variables. Some variables represent possible policy intervention points. A change imposed to one or more of these represents a possible policy intervention. Changes in variables propagate over the causal relations in a time-dispersed way, each relation using its own timeframe. This makes that a change in a variable at the start of a causal relation becomes apparent in the causally dependent variable partly instantly and partly in the longer run. Some relations may pass on just a fraction of the original change, while others may exaggerate the change. Propagation of changes keeps on going between all variables.

To determine the impact of the imposed interventions, some variables are identified as goals. Changes in these variables are measured to find the intervention impact. Some goals may be more relevant than others, even differently for different moments in time.

3.3 Information elements

The information elements included in a MARVEL model are shown in Table 1. New or extended elements as compared to a standard CLD are shown in *italics*.

Table 1 Information elements included in a MARVEL model

Arrow information	**Variable** information	**Relevant moment** information
1. *Name*	1. Name	1. *Name*
2. *Description*	2. *Description*	2. *Time*
3. "From" variable	3. *Minimum value*	3. *Representativeness (different per goal variable)*
4. "To" variable	4. *Maximum value*	
5. *Speed*	5. *Initial value*	
6. *Strength*	6. *Control variable flag*	
7. *Sign*	7. *Control ease*	
8. *Arrow type*	8. *Control lower bound*	
	9. *Control upper bound*	
	10. *Control intervals*	
	11. *Goal variable flag*	
	12. *Goal relevance*	
	13. *Goal direction*	

To **arrows** a *name* and *description* are added to clearly describe and distinguish the causal relations in the model. This information does not only help to validate and document the model. MARVEL supports multiple (parallel) causal relations between two variables and therefore this information is almost indispensable to distinguish these parallel arrows.

Each arrow has a *speed* assigned to it. This is a qualitative term like "fast", "slow" or "mid term". Every arrow also has a *strength* which is a qualitative term like "weak", "1-to-1" or "strong". The strength expresses to what extent a change of the arrow's "from" variable results, in the long run, in a change of the "to" variable. The appropriate terms for arrow speeds and strengths depend on the context of the model at hand.

The traditional CLD arrow *signs* “+” and “-” (or “s” and “o”) are expanded with two new options. One new sign is introduced for more refined behaviour than a standard “+” or “-” sign can offer. It defines in a graphical way how the arrow’s sign varies from fully “+” to fully “-” and anything in between based on the arrow’s “from” and/or “to” variable value. This sign is called the “c” (or “complex”) sign. The second new sign is called the “||” (or “blocked”) sign and means that the arrow is put out of function; a change in its “from” variable does not result in a change of its “to” variable. This is relevant to eliminate the arrow functionally from the model yet recording that the arrow was considered once during the modelling process, but rejected in the end.

The *arrow type* defines how the arrow affects the “to” variable. The standard arrow type is the so-called *influence arrow*. This arrow type passes each change in the arrow’s “from” variable in a time-dispersed way on to the arrow’s “to” variable. In this way the change in the “from” variable becomes apparent in the “to” variable partly instantly and partly in the longer run. MARVEL allows several ways for this time-dispersed propagation of changes to the “to” variable. These ways are called *influence release patterns* and are discussed in more detail in Section 4.2. These patterns are a key element of MARVEL as they define the behavioural ‘intelligence’ of the arrows. Apart from the *influence arrow* two other arrow types are discerned: the *force-value arrow* and the *mapped force-value arrow*. These arrow types do not propagate changes but instead try to force the “to” variable to a value that is dictated by the “from” variable taking into account the arrow’s *speed*[3]. Therefore the *(mapped) force-value arrow* may generate changes in the “to” variable even if the “from” variable does not change.

To **variables** information is added on their role in the model. Some variables represent possible intervention points for the policy to be tested; these are called *control variables*. Other variables represent performance measures to evaluate the effects of a policy; these are called *goal variables*.

For each variable the *minimum* and *maximum value* define the reasonably valid value range for that variable. The *initial value* is the current value of the variable, or more precise the variable’s ‘status quo’ at the moment the policy to be tested is imposed.

For control variables the *control variable flag* is set, marking that they represent an intervention point. The *control ease* specifies how easily a control variable can be changed directly through interventions. It is a qualitative term like “average ease”, “less ease” or “very easy”. This ease is not to be confused with the attractiveness of the control variable as this could also include an appraisal of the expected effects from using the control variable; finding the effects of a control variable is a main purpose of the model itself. The *control lower bound* and *control upper bound* specify the

[3] The *force-value arrow* tries to force the “to” variable to the same value as the “from” variable. The *mapped force-value* tries to force the “to” variable to a value read from a “map” (graph) translating every “from” variable value into the value to which the “to” variable should be forced. All values referred to here are transformed from the variable’s valid value range to the normalized range [0..1]. Several detailed remarks can be stated on the *(mapped) force-value arrow* and its relation to the *complex arrow sign* but these are considered outside the scope of this paper. Details are recorded in an internal design document *(van Zijderveld 2007)*.

boundaries between which a control variable can be actively set to model a policy intervention. The *control intervals* define the number of steps to use between these bounds, to evaluate several intermediate control or intervention levels. The control bounds should lie within the variable's valid value range.

For goal variables the *goal variable flag* is set, indicating that they are relevant to assess the effects of the policy to be tested. The *goal relevance* is a measure expressing the relevance of each goal variable for the overall model behaviour assessment when evaluating the effects of imposed interventions. It is a qualitative term like "average relevant" or "more relevant". The *goal direction* specifies what changes in the goal variable are considered desirable, compared to its initial value. Possibilities are "the greater the better", "the smaller the better" or "no change is best".

Relevant moments are completely new elements compared to a standard CLD. These are the moments in time that are relevant for MARVEL's dynamic behaviour analysis as discussed in Section 4.2. They include at least the start moment for the dynamic behaviour analysis and the furthest time horizon to be covered by the analysis, for instance "now" and "long term". They can be expanded with any moment in time between these extremes if considered relevant reporting moments for the dynamic behaviour analysis, for instance "short term" and "mid term". The *representativeness* is a value expressing the relevance of that moment when establishing in the dynamic behaviour analysis the overall appraisal of a certain combination of interventions. The *representativeness* of a *relevant moment* can be different for different goal variables. This allows specifying that one variable plays a different role in the overall intervention appraisal, say, in the short term than it does in the long term.

3.4 Numeric values behind qualitative terms

Behind the scenes, MARVEL's analysis features are quantitative of nature in spite of the used qualitative terms as will be discussed in Sections 4.1 and 4.2. For this reason there is a numeric value behind each qualitative term for control ease, goal relevance, speed and strength. Also the moments in time that are relevant for dynamic behaviour analysis as discussed in Section 3.3 are specified as numeric values behind the qualitative terms. Proper determination of the numeric values is required to use MARVEL in a sound way as discussed in Section 6.

3.5 Diagram notation

To display the MARVEL specific information elements from Table 1 in a diagram, some extensions to CLD drawing conventions are made. One extension is that the name of an arrow can be displayed next to an arrow as shown in Figure 1. This is especially relevant in case of parallel causal relations.

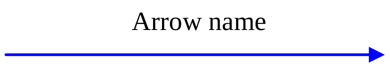

Figure 1 Arrow name display

The arrow's speed is displayed using delay stripes. This is based on the generally accepted delay symbol. The more delay stripes the greater the arrow's delay (or the lower its speed). Each speed term present in a model should have its own number of stripes. Half stripes are also used because the number of speed terms present in a model may be too great to use only whole stripes. At most one half stripe can be used on the arrow, and it must be placed farthest from the arrow head of all stripes. The top row of Figure 2 first shows an arrow with a quite low speed having two and a half stripes; then a somewhat faster one having two stripes, then an even faster one having one and a half stripe and so on. An arrow having no stripes at all is the fastest category.

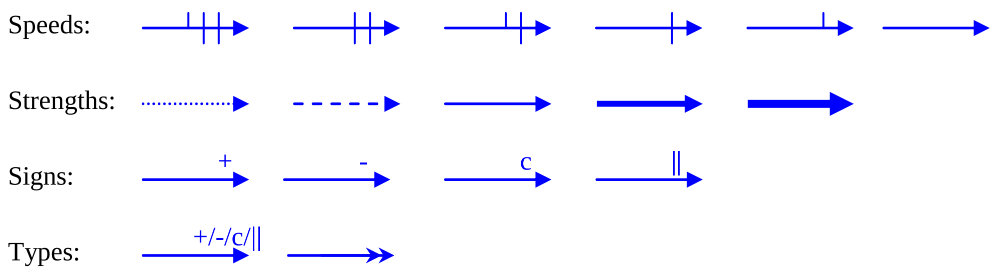

Figure 2 Additional arrow symbols

The arrow's strength is displayed using its line type and thickness. The line features to use are as follows and shown in Figure 2, starting at the weakest category available: dotted line with line thickness 1; dashed line with line thickness 1, solid line with thickness 1; solid line with thickness 2; solid line with thickness 3; and so on. There is no need to have all these categories present in a model.

The arrow's sign is displayed near the arrowhead as shown in Figure 2. Standard CLD signs "+" and "-" sign display as usual while the complex arrow sign displays as "c" and the blocked arrow sign as "||".

The arrow's type becomes apparent through its displayed sign and arrowhead. The arrow sign is only displayed for the *influence* type arrow. For the *force-value arrow* and the *mapped force-value arrow* the sign is not relevant and therefore is not displayed. Instead, they have a double arrowhead as shown at Figure 2's bottom line.

Figure 3 Additional variable symbols

For variables symbols are added that indicate whether it is a control and/or a goal variable. These symbols are shown in Figure 3: a control knob for control variables and a bull's eye for a goal variable. The control knob is placed at the left hand side of the variable's name to stress it is an input to the model, while the bull's eye is located at the right hand side of the variable's name indicating it is an output. A variable can be both a control and a goal variable. This can be useful for instance if a variable is initially changed directly to model an intervention, but also is part of feedback loops further changing its value over time.

4 Analysis with MARVEL

MARVEL includes two approaches to analyse a model: *static loop analysis* and *dynamic behaviour analysis*. Their principles are stated in Section 4.1 and 4.2 respectively[4].

4.1 Static loop analysis

The static loop analysis aims to provide a basic insight into the feedback loops present in the model. It also aims to give a first impression of each feedback loop's power compared to all other loops in the model. This power is based on the feedback loop's total speed and strength as follows.

Two types of power are defined, the *relative power* and the *absolute power*. The latter serves as an intermediate result needed for the calculations[5].

The loop's *absolute power* combines information on the speed, strength and sign of every arrow in that loop. The absolute power is the product of the loop's total speed and strength. The loop's total speed in its turn is constructed such that the loop's duration equals the summed durations for all steps in the loop. The loop's total strength equals the average strength of all steps in the loop thus eliminating any effects of loops seeming to have a great power just because they have a great number of steps. Special measures are taken to include the effects of complex arrow signs that are neither "+" nor "-" and to take account for sets of two or more parallel arrows that may occur at any step in the loop and are likely to differ in speed, strength and sign.

The loop's *relative power* equals the loop's *absolute power* divided by the greatest *absolute power* of any loop in the model, multiplied by 100%. More than one loop in the model may have a 100% relative power.

Results of the static loop analysis are:

1. list of the feedback loops present in the model, and for each feedback loop:
2. total number of steps in the loop;
3. total sign of the loop ("+", "-" or undetermined);
4. relative power of the loop;
5. details on each step in the loop (which arrow, sign, speed and strength);
6. list per variable (or combination of variables) which loops they are part of.

A static loop analysis generally gives less insight than a dynamic behaviour analysis. One reason is that it does not include the interventions that are imposed on control variables. It also does not take into account the room each variable has to change, defined by its valid value range and initial value: a feedback loop having a great power according to the static loop analysis in fact may not function at all if one or more variables in the loop can not change anymore in the desired direction. A last reason is that the static loop analysis only reports on the initial ("current") state as defined in the

[4] Further details not discussed here are recorded in an internal design document *(van Zijderveld 2007)*.
[5] The *absolute power* therefore is not an analysis outcome, while the *relative power* is.

model: it does not report on expected future model states, and therefore not on intervention effects.

4.2 Dynamic behaviour analysis

The dynamic behaviour analysis aims to provide an insight into the model's behaviour over time and in particular in the effects to be expected from interventions on identified control variables. It uses a time stepped approach. In short, MARVEL imposes changes on the control variables according to the *control boundaries* and *control intervals* as specified at each individual control variable. This could either be a combination of several control variables or a single control variable at the time. Each combination of control variables that are actually used plus the values they are controlled to is considered a separate calculation.

Then changes are imposed on these control variables: at the start time for the dynamic analysis they change from their initial value to the control value they should assume for that calculation. These changes propagate through the model over the arrows in smaller time steps, taking into account the speeds, signs, strengths and 'intelligent' behaviour of the arrows. The result is the dynamic behaviour of the model which is reported for the relevant moments (using specific performance measures explained below) and for all time steps (time series showing each variable's value at that time).

MARVEL discerns several performance parameters to measure the performance of specific (combinations of) control variables and the applied control values. These performance parameters are calculated for each *relevant moment* as introduced in Section 3.3.

1. *Value changes*. These can be calculated for all variables and simply are the normalized value of a variable at the relevant moment minus its normalized initial value. For control variables also the initial value is used instead of the applied control value. A value change can have any value in the range [-1..+1] as for dynamic behaviour analysis MARVEL normalizes each variable's valid value range to [0..1].
2. *Gains*. These can only be calculated for goal variables. The gain equals the goal variable's value change times its goal relevance, times a factor expressing the desirability of the value change direction. If the goal variable's goal direction is "the greater the better", then value increments are counted as positive gains and value decrements are counted as negative gains. The opposite holds for goal direction "the smaller the better". For goal direction "no change at all is best" any value change is counted as a negative gain.
3. *Scores*. These combine for one relevant moment information on all goal variables and all active control variables. It is calculated as the sum of the gains of all goal variables for the selected relevant moment, multiplied by the average ease of all control variables that were active during the calculation. The reason to use the average ease is that in MARVEL a calculation where two control variables were active having an "Average ease" is considered equally attractive as a calculation in which only one control variable having "Average ease" was used.

Finally an overall *calculation appraisal* is calculated. This appraisal takes into account that some of the *relevant moments* in the model are more representative for the model's overall performance than other *relevant moments*. This is in accordance with Table 1. Also within one relevant moment, some goal variables may be more representative than others. For instance, one goal variable is mainly representative for the short term performance, while another is more representative for the long term performance. The *calculation appraisal* is determined as follows. Within each relevant moment, for each goal variable its gain at that moment is multiplied by the representativeness of the goal variable within the relevant moment. Then these values are summed over all goal variables and all relevant moments, resulting in the *calculation appraisal*.

Essential in this dynamic behaviour analysis is the new role of the arrows. Consider for example the arrow from Figure 4 defining a negative causal relation between the product unit price and sales[6]. The arrowhead and the presence of an arrow sign tell that it is a standard *influence* arrow that will propagate changes in product unit price to sales.

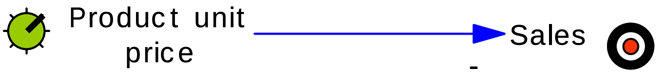

Figure 4 Example arrow

Figure 5 shows the internal construction of such an *influence type* arrow: the gray box equals the single example arrow from Figure 4. The influence released by an *influence* arrow to its "to" variable is based on previous changes in the "from" variable and the part of those previous changes that already have been propagated to the "to" variable. The *influence* arrow accumulates all changes in the "from" variable, corrected for the arrow's sign and strength. From this accumulation the influence is passed on to the "to" variable, taking into account the arrow's speed and the calculation's time step size.

Therefore the influence increment according to Figure 5 during a certain time step in the dynamic behaviour analysis equals the change in value of the "from" variable during that time step, multiplied by the arrow's strength and sign. A "+" sign results in a sign multiplication factor of +1, a sign "-" in a sign multiplication factor of -1, a "||" sign in a sign multiplication factor of 0 (effectively ignoring the value change) and a "c" sign in a value between +1 and -1 according to the to graph defining the sign based on the current "from" and/or "to" variable values.

[6] The product unit price is considered a control variable and the sales a goal variable. This has no effect on the arrow's behaviour.

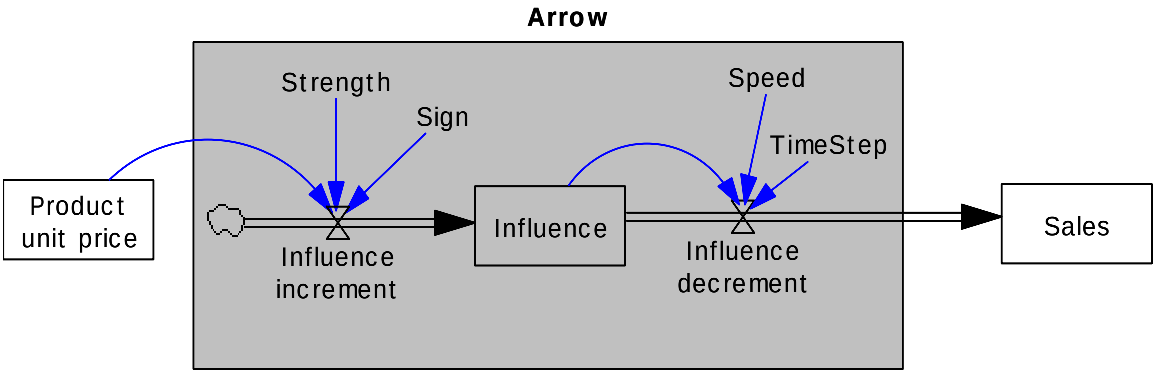

Figure 5 Construction of the MARVEL *influence* type arrow

The influence decrement according to Figure 5 depends on the influence release pattern of the specific arrow. The standard MARVEL influence release pattern is a negative exponential pattern as often used in the field of System Dynamics. This means that the influence to release to the "to" variable equals the size of the influence multiplied by the arrow's speed and the calculation's time step. The same amount is removed from the influence present on the arrow.

An example of this release pattern for the arrow from Figure 4 is shown in Figure 6. In this example goal variable "Sales" has an initial value of 0. Control variable "Product unit price" has an initial value of 1 and is controlled to a value of 0 halfway the short term timeframe. The *influence type* arrow has a negative sign. Some choices on the arrow speed and strength are made[7].

Figure 6 shows the essence of MARVEL: a single onetime change in the arrow's "from" variable results in a prolonged effect on the "to" variable, using the arrow's influence release pattern.

[7] I.e. default settings of the Tool Implementing MARVEL (TIM) as introduced briefly in Section 5.

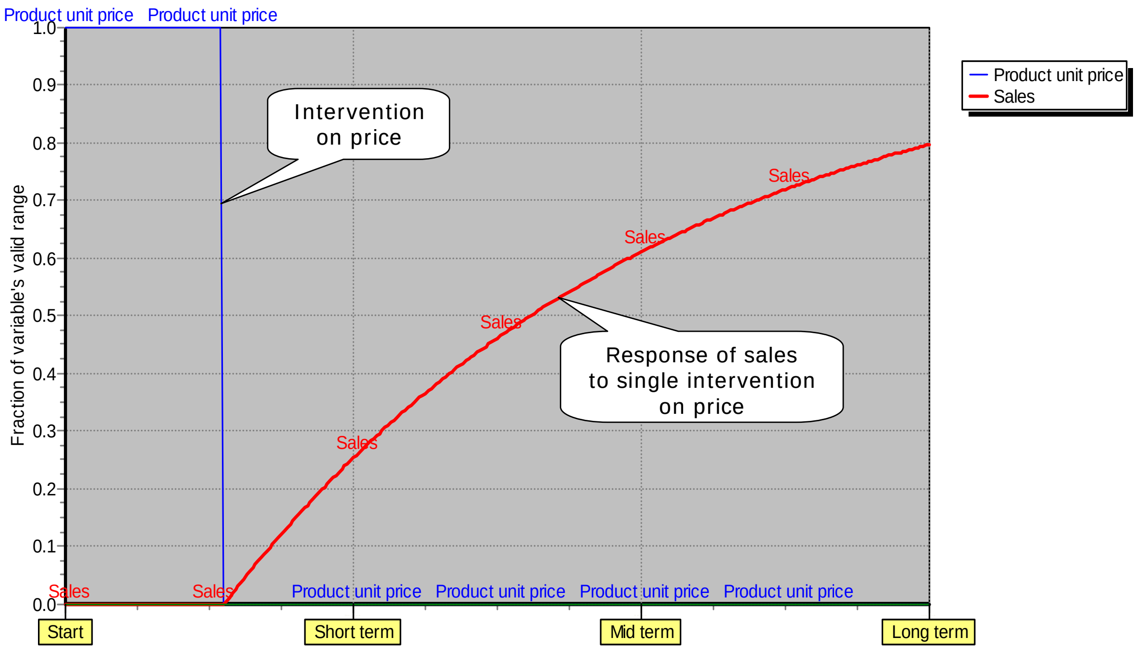

Figure 6 Standard (negative exponential) influence release pattern; prolonged effect on sales from a single change in product unit price (from 1 to 0)

5 Illustrative example

To give a brief impression of a MARVEL diagram Figure 7 shows a fictitious case. It is inspired on the project from which the MARVEL approach originated. In short, this problem is about finding ways to speed up the introduction of a new software package in a company, having great impact on the procedures and organisation. Resistance against the new software hampers its introduction, caused by potential future dismissals when the software is introduced, by low software quality and poor management.

Four possible intervention points are identified, each having its own range to be controlled to:

1. Available budget;
2. Clear mandate assignment;
3. Management knowledge level;
4. Management quality

These intervention points can be used either individually or in any combination. To every causal relation a speed, strength and sign is assigned as displayed in the diagram. The actual problem was analysed using MARVEL and turned out to result in accepted and valuable results that are used in practice.

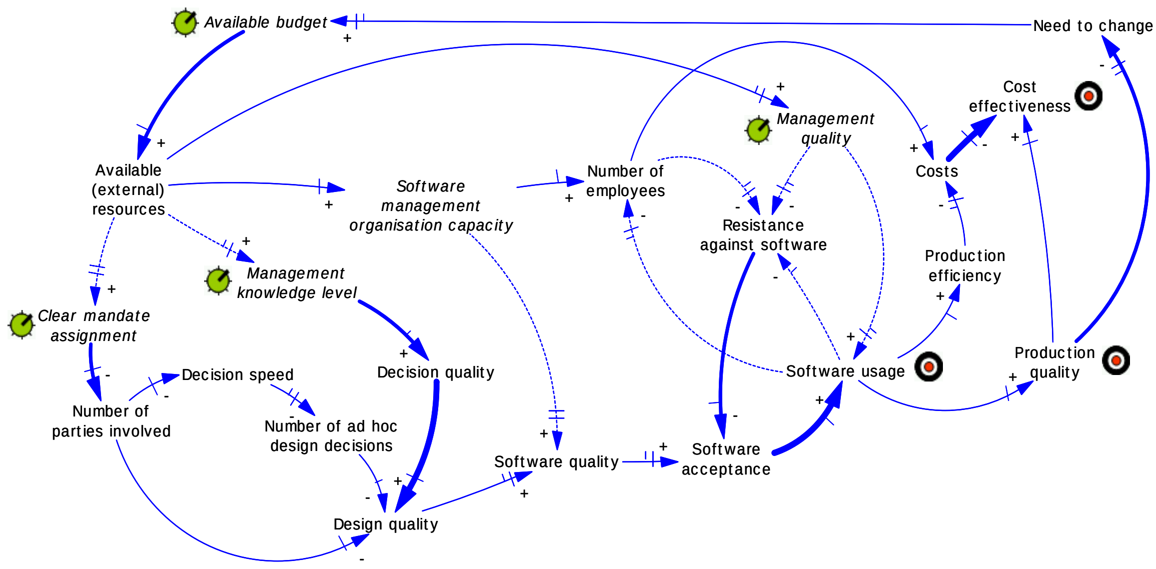

Figure 7 MARVEL diagram for a fictitious case

A full list of the model variables including their use as control or goal variable is given in Figure 8. The screenshot is taken from the Tool Implementing MARVEL (TIM) that was initiated in parallel to the development of MARVEL[8]. For control variables their control ease plus control upper bound and lower bound are depicted. For goal variables their goal relevance and goal direction are visible. In this fictitious case all variables have a valid range [0..1].

[8] At the moment this paper is written TIM is not yet available to third parties. If, however, sufficient parties express to the author their interest in using TIM it may be decided to make TIM available.

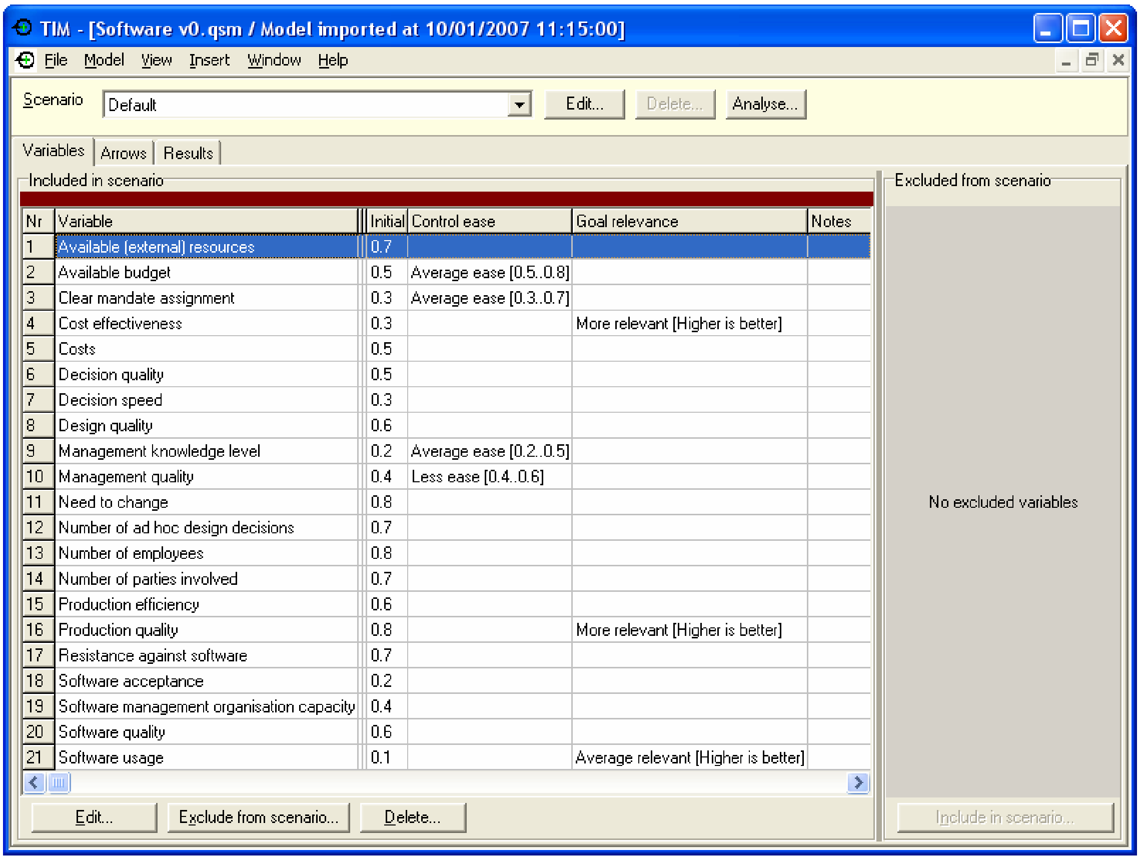

Nr	Variable	Initial	Control ease	Goal relevance	Notes
1	Available (external) resources	0.7			
2	Available budget	0.5	Average ease [0.5..0.8]		
3	Clear mandate assignment	0.3	Average ease [0.3..0.7]		
4	Cost effectiveness	0.3		More relevant [Higher is better]	
5	Costs	0.5			
6	Decision quality	0.5			
7	Decision speed	0.3			
8	Design quality	0.6			
9	Management knowledge level	0.2	Average ease [0.2..0.5]		
10	Management quality	0.4	Less ease [0.4..0.6]		
11	Need to change	0.8			
12	Number of ad hoc design decisions	0.7			
13	Number of employees	0.8			
14	Number of parties involved	0.7			
15	Production efficiency	0.6			
16	Production quality	0.8		More relevant [Higher is better]	
17	Resistance against software	0.7			
18	Software acceptance	0.2			
19	Software management organisation capacity	0.4			
20	Software quality	0.6			
21	Software usage	0.1		Average relevant [Higher is better]	

Figure 8 Screenshot on variables defined for a fictitious case

Some arrows in the model are shown in Figure 9 including their sign, speed and strength. All arrows are defined as standard *influence type* arrows. The qualitative terms used for arrow speed and strength as well as those for control ease and goal relevance can be defined at will in a model including the numeric values behind them.

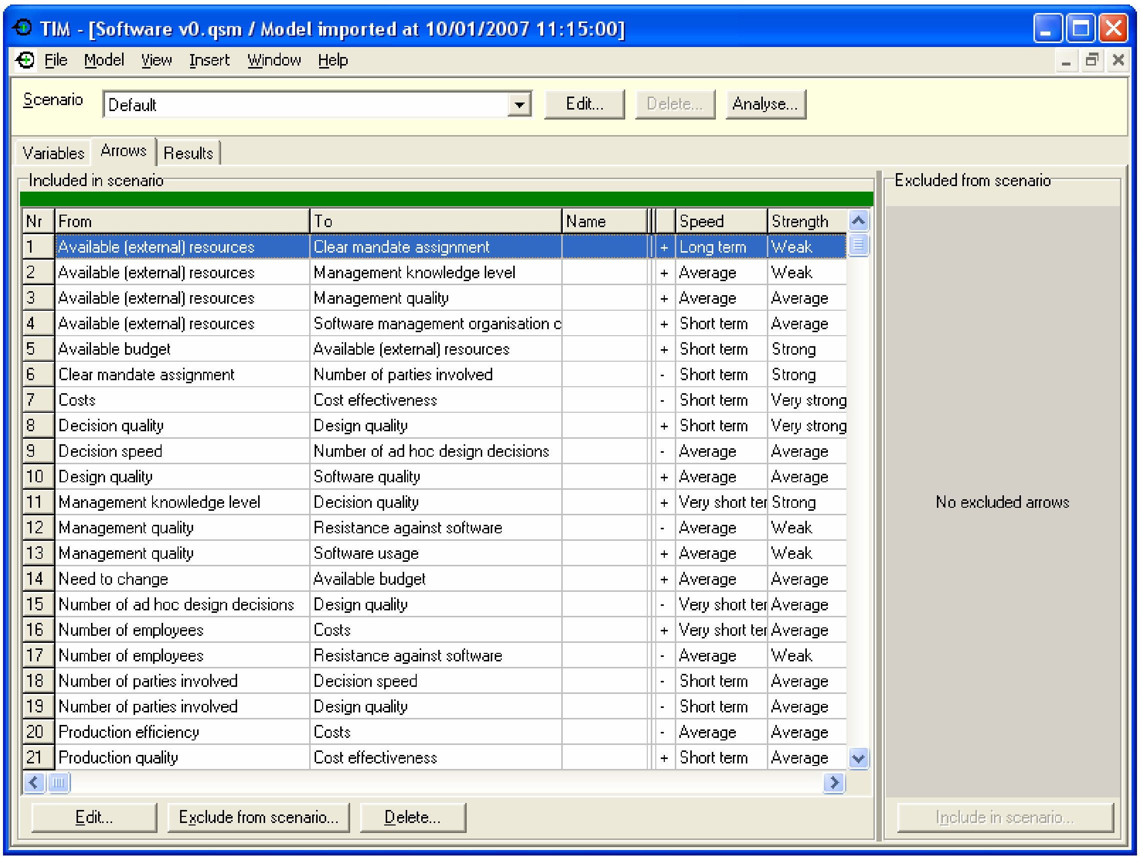

Nr	From	To	Name		Speed	Strength
1	Available (external) resources	Clear mandate assignment		+	Long term	Weak
2	Available (external) resources	Management knowledge level		+	Average	Weak
3	Available (external) resources	Management quality		+	Average	Average
4	Available (external) resources	Software management organisation c		+	Short term	Average
5	Available budget	Available (external) resources		+	Short term	Strong
6	Clear mandate assignment	Number of parties involved		-	Short term	Strong
7	Costs	Cost effectiveness		-	Short term	Very strong
8	Decision quality	Design quality		+	Short term	Very strong
9	Decision speed	Number of ad hoc design decisions		-	Average	Average
10	Design quality	Software quality		+	Average	Average
11	Management knowledge level	Decision quality		+	Very short ter	Strong
12	Management quality	Resistance against software		-	Average	Weak
13	Management quality	Software usage		+	Average	Weak
14	Need to change	Available budget		+	Average	Average
15	Number of ad hoc design decisions	Design quality		-	Very short ter	Average
16	Number of employees	Costs		+	Very short ter	Average
17	Number of employees	Resistance against software		-	Average	Weak
18	Number of parties involved	Decision speed		-	Short term	Average
19	Number of parties involved	Design quality		-	Short term	Average
20	Production efficiency	Costs		-	Average	Average
21	Production quality	Cost effectiveness		+	Short term	Average

Figure 9 Screenshot on arrows defined for a fictitious case

An impression of static loop analysis as discussed in Section 4.1 for the fictitious case of Figure 7 is given in Figure 10. It shows how nine feedback loops are found, one of which has a clearly far greater power than the other ones. That loop contains just fast causal relations: increased software acceptance results in a strong and fast software usage increase; causing a fast yet weak reduction of the resistance against the software; further causing a fast and strong software acceptance increment.

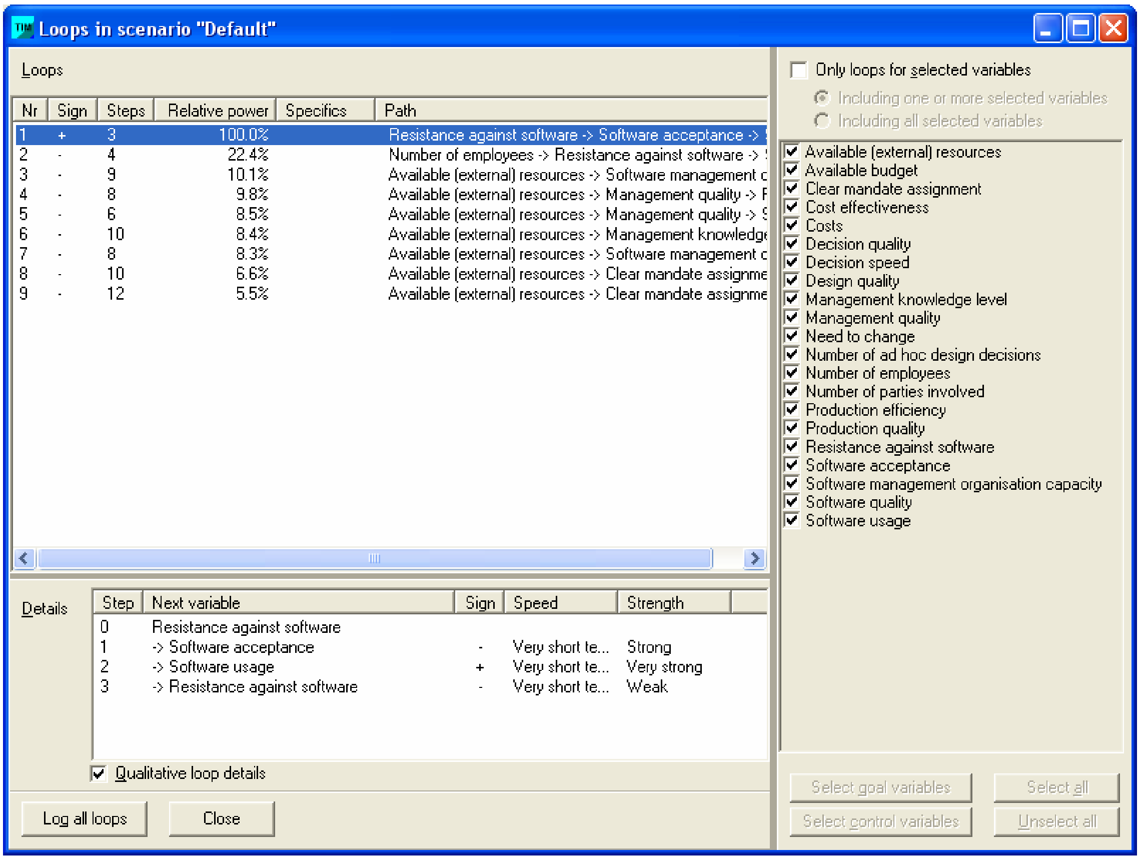

Figure 10 **Screenshot on static loop analysis for a fictitious case**

Figure 11 gives an impression of the dynamic behaviour as discussed in Section 4.2. It shows the response of all control and goal variables if the available budget is increased from its initial level of 0.5 (see Figure 8) to its control upper bound of 0.8. According to Figure 7 the additional budget is transferred into more resources meaning that external people are hired. This initiates changes in clear mandate assignment, management knowledge level, management quality and software management capacity (the latter not shown in Figure 11). Software usage improves because of these changes, either directly (management quality) or indirectly. Increased software usage starts off the fast and strong positive feedback loop: resistance against software is reduced, software acceptance is increased and software usage is further boosted.

Costs initially rise because of an increased number of employees in the software management organisation. When the software is used more and more, the number of employees is reduced because of increased efficiency, thus reducing the costs and increasing the production quality. This causes a rise in the cost effectiveness after an initial dip. These dynamics are all the result of a single change in the available budget.

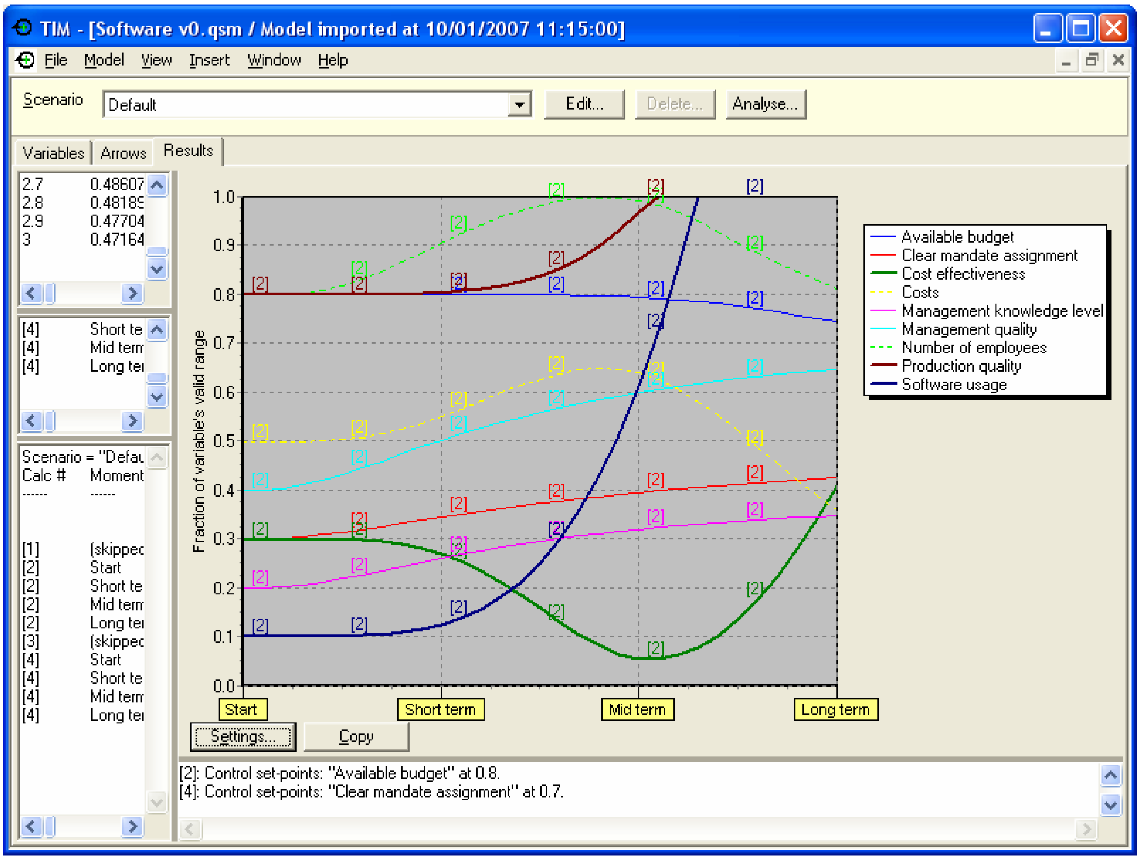

Figure 11 Screenshot with output (time series) after dynamic behaviour analysis for a fictitious case

Figure 12 shows an example of one of the performance measures as discussed in Section 4.2 (the *score*) for the cases each of the control variables from Figure 7 are applied individually plus some combinations. It shows to what extent these possible interventions result in improvements in all goal variables taking into account the ease of the applied control variables. In this example each control variable is applied to its respective *control upper bound* as defined in the fictitious case and shown in Figure 8.

Figure 12 reports that for the mid term good results are found by improving both management quality and management knowledge. Adding budget reduces the overall score for the mid term mainly as improvements in cost effectiveness lag due to initial hiring of external resources thus increasing costs. Changing only management knowledge level has the highest score on the long term. This is caused by not intervening on management quality which is rated as less easy as shown in Figure 8.

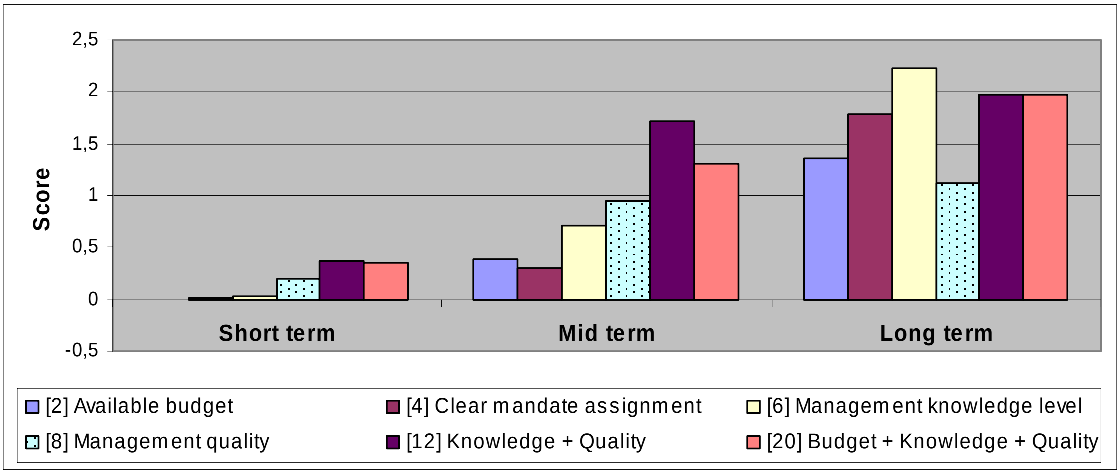

Figure 12 Score of several control variable combinations for a fictitious case

6 Practical and methodological issues

There are several remarks on methodological and practical issues to be made.

First, MARVEL is designed for a first policy intervention analysis for problems with limited quantitative data availability. The advantage of having a method that can be used under these conditions is paid off against limited opportunities to tailor model behaviour as compared to a quantified SFD. This limitation will generally result in less accurate results than can be obtained from an SFD if the quantitative data were available. A MARVEL user should always bear this restriction in mind.

Second, MARVEL focuses on changes. This limits its application to problems where the study of the propagation of changes between model variables is an adequate modelling approach. Flows are not explicitly modelled. Therefore MARVEL can not be used to study the effects of changing inflows or outflows on a stock (or "level") as can be done with an SFD. An alternative approach that may be acceptable in several cases is studying the changing imbalance between inflow and outflow.

Third, MARVEL assumes an initial equilibrium status of the model that is disturbed by an intervention. It is not meant to study the autonomous behaviour of a model without interventions although the (mapped) force-value arrows offer some, but still limited, possibilities for this type of study.

Fourth, the qualitative terms used for speeds and strengths require attention. They often will be problem-specific. The qualitative terms for speeds and strengths plus assigned numeric values must therefore be defined and validated in coherence with the identified relevant moments for the problem. This could be embedded in the GMB process as follows. First design the CLD without MARVEL-specific information elements. Then establish the valid value range for each variable[9]. Subsequently the relevant moments

[9] For some variables the maximum may be hard to find, for instance for "Number of employees". Yet often a reasonable estimate of a realistic maximum can be made. Otherwise a 0-100% scale may be used where 0% means fully absent and 100% fully present or optimal, for instance for "Design quality".

are defined by graphically placing them on a timescale. Then the terms plus numeric values for speeds are defined while graphically showing the response at each identified speed to a standard exemplary intervention. This response must be shown against the background of the identified relevant moments. The standard intervention could be a change in a variable at 50% of its valid value range while using a neutral strength, i.e. strength 1. An analogous approach can be used to define the strengths while choosing a standard speed. Then set the speeds and strengths for the arrows against the background of the valid value ranges of the related variables. If required at this point, a valid value range for a variable can be adapted or the set of speeds or strengths can be extended.

Fifth, the currently defined standard influence release pattern is a negative exponential pattern as often used in the field of System Dynamics. Further research into the causal relation classes for which this pattern is valid and the development of alternative release patterns plus their associated causal relation classes is required.

Finally the question of the added value of MARVEL as compared to a quantified SFD springs to mind. The main part of the answer is found in the starting point that MARVEL is designed for a first policy intervention analysis for problems with limited quantitative data availability. This means that it is intended for problems for which a full SFD model is not feasible or desirable. Another part of the answer is that even if quantitative data lacks, decision makers often can reach consensus on the speed and strength of causal relations. Although limited in accuracy, this information is valuable for an analysis. A quantified SFD does not offer a standard and easy-to-use structure to benefit from this ‘soft’ information, but MARVEL does. A first comparison of MARVEL and SFD indicates that the number of variables in a MARVEL model tends to be around 30% of the number of variables in an SFD. Comparing the results of published models for both methods appears to be difficult because a case with both a consistent and well-documented CLD and SFD is rarely found.

7 Concluding remarks

The principles for a new analysis method called MARVEL to develop and assess potential interventions based on a Causal Loop Diagram are presented. MARVEL includes several new features such as the addition of ‘intelligent’ behaviour to arrows allowing changes in model variables to propagate in a time-dispersed way. Also an integrated approach for intervention performance measurement streamlines the use of this method. Model complexity is scalable because of the inclusion of several optional more advanced features like complex arrow signs or *force-value arrows*.

MARVEL allows fast model development and policy analysis when implemented in a tool like TIM. Proper use of MARVEL however does not escape from reference modes and/or other model validation approaches. Finding the right qualitative terms for speeds, strengths and relevant moments requires attention, and even more do the numeric values behind them. Also setting variables’ valid ranges, selecting the right influence release pattern or constructing a complex arrow sign must be done with care. One should always be aware of MARVEL’s scope and approach of using simplified and generalized behaviour.

Nevertheless, based on experiences with the method in the real-life project during which it was developed, it is believed that this method offers an integrated framework that can be a valuable addition to the field of System Dynamics for ‘soft’ policy type evaluation problems.

Focal points in future developments of MARVEL will include alternative ways to impose interventions, further *influence release patterns* and finding optimal interventions. Also synergetic effects or constraints of multiple causes of a model variable will deserve attention.

References

Coyle, R. Geoff. 2000. *Qualitative and quantitative modelling in system dynamics: some research questions*. System Dynamics Review 16(3): 225-244.

Coyle, R. Geoff. 2001. *Rejoinder to Homer and Oliva*. System Dynamics Review 17(4): 357-363.

Homer, Jack B, Oliva Rogelio. 2001. *Maps and models in system dynamics: a response to Coyle*. System Dynamics Review 17(4): 347-355.

Liddell, Will G., Powell, John H. 2004. *Agreeing access policy in a general medical practice: a case study using QPID*. System Dynamics Review 20(1): 49-73.

McLucas, Alan Charles. 2001. *An investigation into the integration of qualitative and quantitative techniques for addressing systematic complexity in the context of organisational strategic decision-making*. Canbarra: University of NSW.

Sterman, John D. 2000. *Business Dynamics – Systems Thinking and Modeling for a Complex World*. Boston: Irwin McGraw-Hill.

van Zijderveld, Erik J.A. 2007. *MARVEL design document* (internal document).

Vennix, Jac A.M. 1996. *Group Model Building – Facilitating Team Learning Using System Dynamics*. Chichester: John Wiley & Sons.

Wolstenholme, E.F., Coyle R. Geoff. 1983. *The Development of System Dynamics as a Methodology for System Description and Qualitative Analysis*. The Journal of the Operational Research Society 34(7): 569-581.

Abbreviations

CLD	Causal Loop Diagram
GMB	Group Model Building
MARVEL	Method to Analyse Relations between Variables using Enriched Loops
QPID	Qualitative Politicised Influence Diagram
SFD	Stock and Flow Diagram
TIM	Tool Implementing MARVEL