MAIN ARTICLE

Refinements on eigenvalue elasticity analysis: interpretation of parameter elasticities

Sergey Naumov[a] and Rogelio Oliva[b]*

Abstract

The purpose of this article is to report on improvements on the interpretation and insights emerging from *dynamic decomposition weight analysis* (DDWA). These improvements emerged from efforts to further automate and expand the eigenvalue elasticity analysis methods and resolve inconsistencies in assumptions made in published reports of DDWA usage. In addition to making available to the broad system dynamics community an improved toolset to perform eigenvalue elasticity analysis, in this paper we clarify the set of assumptions needed to obtain reliable results and develop a new framework to assess the impact of model parameters on the projections of behavior modes on stock behavior. We illustrate the use of these developments by updating a previously published model analysis. The paper concludes by summarizing our findings and their implications for the further development of structural dominance analysis.

Syst. Dyn. Rev. (2018)

Introduction

The purpose of this article is to report on improvements on the interpretation and insights emerging from *dynamic decomposition weight analysis* (DDWA) originally developed by Saleh *et al.* (2010). These improvements emerged from efforts to further automate and expand the eigenvalue elasticity analysis (EEA) methods and resolve inconsistencies in assumptions made in published reports of DDWA usage. It should be noted that these improvements and inconsistencies could only be detected when attempting to expand the domain of EEA established in papers developing the proof-of-concept for the method. We expect that further improvements will emerge as we expand the use of the tools to different types of models and with different goals. It is in this spirit that we make the toolset available to the system dynamics community (Naumov and Oliva, 2018).

Oliva defines the purpose and workings of EEA as follows:

> Eigenvalue elasticity analysis (EEA) is a set of methods to assess the effect of structure on behavior in dynamic models (Kampmann, 2012; Kampmann and Oliva, 2006; Oliva, 2015). It works by considering model behavior as a

[a] System Dynamics Group, MIT Sloan School of Management, 100 Main Street, Cambridge, Massachusetts 02142, U.S.A.
[b] Mays Business School, Texas A&M University, TAMU 4217, College Station, Texas 77843, U.S.A.
* Correspondence to: Rogelio Oliva. E-mail: roliva@tamu.edu

Accepted by James Duggan, Received 2 May 2018; Revised 21 May 2018 and 25 July 2018; Accepted 1 August 2018

System Dynamics Review
System Dynamics Review
Published online in Wiley Online Library
(wileyonlinelibrary.com) DOI: 10.1002/sdr.1605

> combination of characteristic behavior modes and assessing the relative importance of particular elements of system structure in influencing these behavior modes. … EEA uses linear systems theory to (i) decompose the observed behavior into its constituent behavior modes, such as oscillation, growth and exponential adjustment; and (ii) outline how a particular behavior mode and its appearance in a given system variable depend upon particular parameters and structural elements (links and loops) in the system. (Oliva, 2016, pp. 26–27).

The main advantage of EEA relative to other experimental methods used for structural dominance analysis is that it provides an exhaustive and precise account of the relationship between model structure and model behavior (Duggan and Oliva, 2013). Over the years, EEA methods have evolved as new developments become available (for an overview of historical developments, see Kampmann and Oliva, 2008, p. 511; Oliva, 2015, pp. 209–210). The current version of the EEA process is summarized in Figure 1. After the original model is linearized, it is possible to calculate the eigenvalues and eigenvectors of the system matrix. Structural explanations of the system behavior are obtained by assessing the influence of the model's feedback loops on the behavior modes; this is referred to as *loop eigenvalue elasticity analysis* (LEEA). The identification of high leverage policies is done by assessing the effect of model parameters on the projections of the behavior modes in the state variable of interest, and this is referred to as *dynamic decomposition weight analysis* (DDWA).

To date, the most comprehensive usage of the EEA is documented in Oliva (2016). In that paper, Oliva explored the usefulness of the toolset in a more realistic model than the simple and stable models that up to that point had been used to test and develop the toolset. Specifically, he expanded the application domain of these methods by increasing the model size and incorporating stochastic variance in some model variables. As far as we can tell, the process and results reported in pages 43–47 of the 2016 paper are an accurate description of the analysis performed at the time and a correct description and interpretation of the policy analysis performed by Oliva and

Fig. 1. Schematic representation of EEA process.

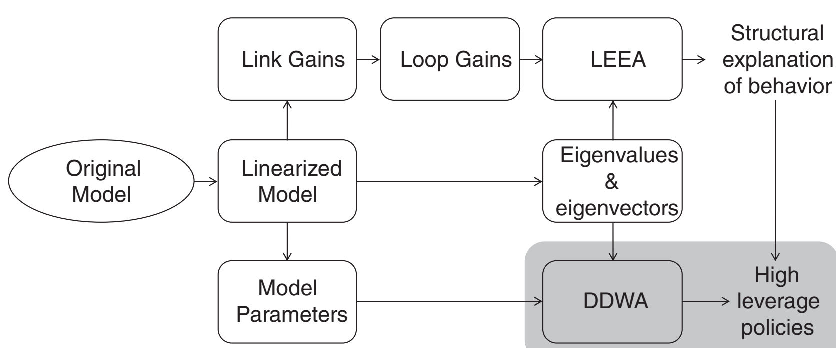

DOI: 10.1002/sdr

Sterman (2001) and the results they obtained. However, further exploration of the tools for EEA suggests that a different set of assumptions to perform the structural decomposition weight analysis reveals better insights. Specifically, in the 2016 paper Oliva performed the evaluation of the weight elasticity to parameters by including in the analysis the effect that parameters had on the model *initial conditions.* That is, when assessing the effect of a parameter on the dynamic decomposition weights, the evaluation (table 3 in the paper) also included the effect that parameters had on initializing state variables. This strategy was selected as it provided a better mapping of the sensitivity analyses performed by Oliva and Sterman and had the additional advantage of allowing the model to be used as is: fewer changes to the model structure—a major concern when assessing the efficiency of the toolset.

This assumption, however, turns out to be problematic for policy analysis using dynamic decomposition weights for three reasons:

1. It gives a false sense of parameter influence if the parameters are used to initialize the model stocks—a common practice to initialize model sectors in equilibrium.
2. Changes in initial conditions might result in unanticipated transient behaviors (Moxnes and Davidsen, 2016).
3. Model initial conditions are seldom effective policy levers as they represent the current (or past) state of the system and policymakers need to determine the course of action given the current state of the system.

Indeed, while initializing a model in equilibrium in terms of other variables and parameters is convenient for model testing (see Sterman, 2000, §18.1.5), one must realize that this is a modeling strategy for *testing* expedience. For policy analysis purposes, initial conditions should be treated as model constants that are different from those parameters that have an influence throughout the duration of the simulation. Note that the effect of this assumption on DDWA had not been empirically explored before because, as described in Oliva (2016), up to that point all EEA applications were on small models built explicitly for demonstration purposes.

In the latest implementation of EEA—which we call the structural dominance analysis (SDA) toolset (Naumov and Oliva, 2018)—we have corrected this assumption and now weight elasticities to model parameter are computed *after* the model has been numerically initialized. This new assumption also enables the computation of these elasticities to be performed at any point during the simulation and not only at the initial simulation time—the only option available with the original implementation of DDWA. The new ability to perform DDWA at any point in the simulation, however, forced us to develop a new interpretation of the impact of model parameters on model behavior and policy development. Note that this change of assumptions only

DOI: 10.1002/sdr

affects the computations of DDWA and the identification of high-leverage policies (shaded area in Figure 1). The rest of the EEA, i.e. model linearization, computation of eigenvalues and eigenvectors and LEEA, remains unaltered.

In the following sections, we provide a new framework to assess the impact of model parameters on the projections of reference modes on stock behavior under the new set of assumptions, and we update and reinterpret the analysis of the Oliva and Sterman (2001) model under these new assumptions. We conclude by summarizing our findings and their implications for the further development of structural dominance analysis.

Identifying influential parameters

In the absence of changes in exogenous inputs, the resulting behavior of any given state variable (t) of a linear system can be written as a weighted sum of a set of behavior modes:

$$x_i(t) = w_{i,0} + w_{i,1}e^{\lambda_1 t} + \ldots + w_{i,n}e^{\lambda_n t}$$

where the λs are the eigenvalues of the Jacobian matrix of the linearized system and the weights *w*s are constants that depend upon the eigenvectors and the initial conditions of the system (see Saleh *et al.*, 2010, for derivation). Each of the system eigenvalues represents a behavior mode. For real eigenvalues, the behavior mode in the form $e^{\lambda t}$ describes exponential growth ($\lambda > 0$) or decay ($\lambda < 0$). Complex eigenvalues appear in conjugate pairs $\lambda \pm i\omega$, which give rise to behavior modes in the form $e^{\lambda t}\sin(\omega t + \theta)$ that describe oscillations with frequency ω within the envelope $e^{\lambda t}$ that is either expanding ($\lambda > 0$) or damped ($\lambda < 0$). In its original conception, DDWA pointed to policy recommendations by isolating the system parameters that affect the projection or weight (w) of a particular *behavior mode* on a stock of interest of the linearized system (Saleh *et al.*, 2010; Oliva, 2015). We have found, however, that it is more useful to assess the influence of a parameter change on the *envelope* of the behavior modes, i.e. the real part of the system eigenvalues.

Let us consider the projection of a behavior mode determined by the real part of an eigenvalue:

$$r(t) = we^{\lambda t} \tag{1}$$

For real eigenvalues, Eq. (1) is the complete behavior mode; for complex eigenvalues, this is half of the envelope of oscillations, i.e. the upper bound on possible values. Note that EEA is performed on a system that has been

DOI: 10.1002/sdr

Fig. 2. Behavior mode convergence rate.

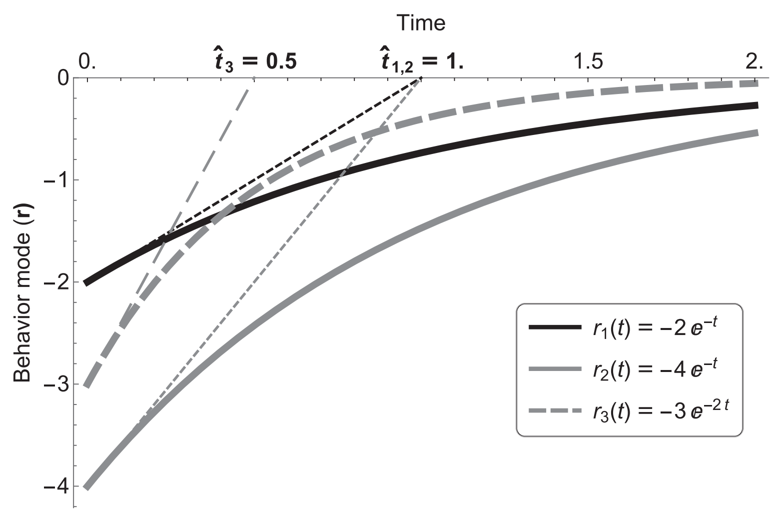

linearized around a specific operating point at model time T, which in the new version of EEA can be any time over the model simulation horizon. From a perspective of the DDWA, this model time T becomes the starting time for the analysis, or $t = 0$. In Eq. (1), the weight w defines the initial value of the projection or the envelope of the behavior mode (at $t = 0$, the equation becomes simply $r(0) = w$), while λ influences the rate of convergence of the exponential trajectory over time. The latter can be evaluated as the convergence time $\hat{t}$, where the instantaneous time derivative $\dot{r}(t) = w\lambda e^{\lambda t}$ reaches zero (see Figure 2). At $t = 0$, this derivative becomes $\dot{r}(0) = w\lambda$, so we can solve for $\hat{t}$:

$$-w + w\lambda\hat{t} = 0 \Rightarrow \hat{t} = \frac{1}{\lambda}$$

Note that the weight w shifts the envelope of the behavior mode up or down (an instantaneous effect), while λ is responsible for the convergence or divergence of the envelope of the behavior mode (an effect that is manifested through time). For instance, in Figure 2, the behavior mode r_3 with the smallest $\hat{t}$ (and the largest λ) shows the quickest convergence, despite the fact that it has a higher initial offset (w) than the behavior mode r_1. By extension, the parameter that results in the largest change in λ for the behavior mode of interest is the most effective long-term policy lever, but the short-term effect of the parameter in w also needs to be considered.

Figure 3 shows the change in the envelope of the behavior mode depending on the sign of w and λ, and the relative magnitude of the weight and eigenvalue elasticities to parameters (for complex eigenvalues, this represents only half of the envelope). The top left quadrant baseline (bold line)

DOI: 10.1002/sdr

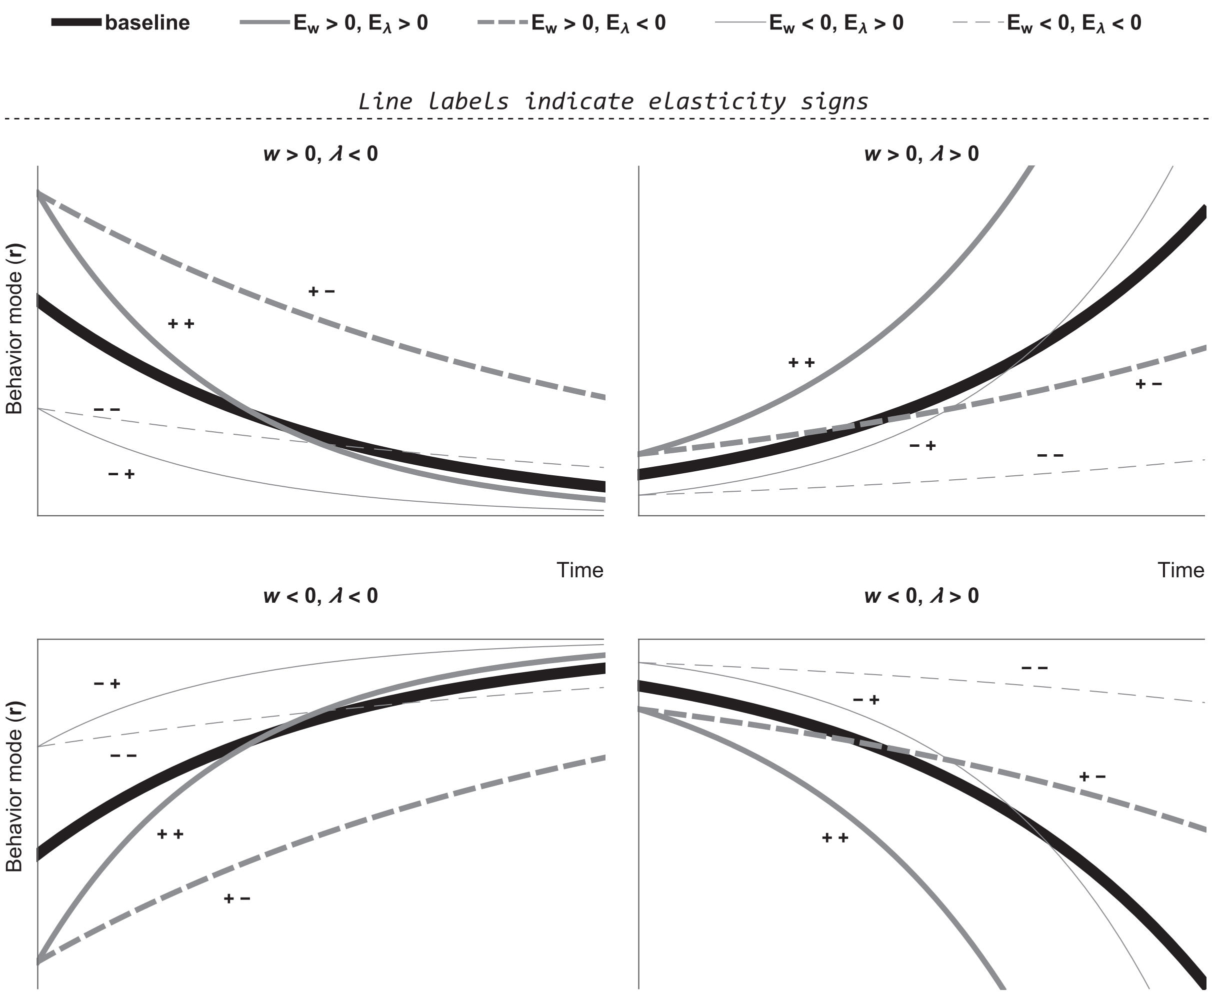

Fig. 3. Weight and behavior mode elasticity to parameters.

shows the envelope of an exponential decay reference mode ($\lambda < 0$) with a positive projection ($w > 0$) on the stock. The solid dark-gray line (labeled **++**) shows the effect of augmenting a parameter whose weight and eigenvalue elasticities are positive ($E_w > 0$; $E_\lambda > 0$). Such change would result in a higher starting point for the projection of the reference mode (the positive elasticity of a positive weight w means a larger value at $t = 0$ in Eq. (1)), but a faster convergence rate (the positive elasticity of a negative eigenvalue λ means more negative value of the exponent in Eq. (1). The light-gray dashed line (labeled **−−**) shows a lower starting point for the projection of the reference mode and a slower convergence rate that would emerge from

DOI: 10.1002/sdr

augmenting a parameter whose weight and eigenvalue elasticities are negative ($E_w < 0$; $E_\lambda < 0$). The other two lines (**+−** and **−+**) show the other two possible combinations of the elasticity signs. The identification of the most desirable parameters for policy design is determined by the desirability of the change (does one want to increase or decrease the projection of the eigenvalue) and the effect of the parameter on the behavior mode. The rest of the quadrants in Figure 3 show the different sign combinations of w, λ, E_w, and E_λ.

We believe that this framework provides an intuitive way to capture the dual effect of a parameter on the behavior mode and its projection on a stock. Accordingly, the updated SDA toolset provides a table with all the required information to identify the model parameters that represent the highest leverage point to augment or diminish a particular behavior. The next section shows the interpretation of this improved output for the same analysis performed in the DDWA section of Oliva (2016).

Analysis of Oliva and Sterman (2001)

Oliva (2016) focused its analysis on the policy goals articulated in section 5 of Oliva and Sterman (2001), i.e. how to reduce the erosion of service quality (see Appendix for an overview of the model structure). Table 1 reports the weight and eigenvalue elasticities of eigenvalue 13 (the long-term decay) on *desired time per order* (the model's proxy for service quality) at time 53 (the simulation time when the DDWA was performed) for the model analyzed by Oliva (2016). Note that at time 53 the real part of eigenvalue 13 is negative ($\lambda_{13} = -0.00083$) and the weight on stock *desired time per order* is positive ($w_{13} = 1.79$); thus the projected behavior mode falls into the upper left quadrant of Figure 3 (see also figure 9 in Oliva, 2016). The table only reports the top 13 parameters, as the rest have elasticities three orders of magnitude smaller than the smallest reported value. This table is the equivalent of table 3 in Oliva (2016), with two important differences. First, the reported elasticities in this table are computed under the new assumption; i.e. the model is initialized numerically using the original values of all parameters, so any change in parameter values during the analysis does not affect the initialization. It should be noted, however, that in this model only two (out of 13) stocks (*service backlog* and *desired labor*) were initialized in equilibrium using other variables and parameters. The second change is that in this table, as per the insights developed in the previous section, the parameters are ranked according to the elasticity of the real part of the eigenvalue, rather than the elasticity of the weight.[1] As a reference, the elasticities of the

[1]The toolset allows for different sorting of the parameters, e.g. alphabetical and by elasticity of weight and real and imaginary parts of eigenvalues.

DOI: 10.1002/sdr

Table 1. Parameter elasticity of weight (w) and behavior mode (λ) of eigenvalue 13 on desired time per order

Rank	Parameter	Elasticity of w of λ_{13} on desired time per order	Elasticity of $\mathrm{Re}(\lambda_{13})$	Elasticity of w of λ_{13} on desired time per order (from Oliva, 2016)
1	Hours per week per employee	−9.622	−3.733	−23.261
2	Desired delivery delay	−9.622	−3.733	0.000*
3	Rookie effectiveness	0.951	−1.458	−0.004*
4	Time to adjust DTO down	−0.161	−0.553	0.000*
5	Alpha	0.142	0.516	0.356
6	Time to adjust desired labor	0.099	−0.288	0.099
7	Fraction of personnel for training	−0.136	0.208	0.001*
8	Time to adjust labor	0.055	−0.180	−1.259
9	Beta	−0.002	−0.037	−0.001*
10	Time for turnover	0.002	0.019	−0.008*
11	Time to perceive labor productivity	0.002	0.014	−0.533
12	Time for experience	0.001	−0.010	0.062
13	Hiring delay	−0.000	−0.001	0.667
	…			

*Not reported in Oliva (2016).

weight originally reported by Oliva (2016) are listed in the last column of Table 1.

When comparing the contents of the table above to table 3 in Oliva (2016) several differences are noteworthy. First, all the initialization parameters that were present in the original table have been dropped from the updated table. This is consistent with the purpose of the analysis—indeed, these parameters were ignored in the discussion in the 2016 paper. Second, the weight elasticities to constants used in stock initializations are significantly modified. The weight elasticity to *hours per week per employee* has been cut by almost 60 percent and *desired delivery delay* has now moved to be one of top-ranking policy levers. By removing the effect of initial conditions from the elasticity computations, we now see in the updated table the long-term effect of parameter changes on the desired policy intervention. This, we believe, is a clearer interpretation of the effect of parameter changes in the policy design. Indeed, increasing the *desired delivery delay* is another policy that was not explored by Oliva and Sterman and would be highly influential as it would remove much of the *work pressure* observed in the system—the driver of the “cutting corners” behavior.

The third significant change is in the elasticity of *time to adjust labor.* This is due to the fact that its effect on the equilibrium initialization of the *desired labor* stock is no longer captured in the analysis. The main effect of the change of assumptions is the change of sign of the elasticity (from negative in the original table to positive here). While the analysis presented in Oliva (2016) still

DOI: 10.1002/sdr

explains why the faster capacity acquisition policy (policy 1) was ineffectual in the original tests performed by Oliva and Sterman (2001)—when changing parameter values they were also changing the model initial conditions—the updated results indicate that their intuition was correct and that *time to adjust labor* and the *hiring delay* should move in the same direction as assessed by the elasticity of the eigenvalue to changes in these parameters.

Finally, the removal of the effects of the initial conditions increased the rank of *alpha* and *time to adjust DTO down* on the projection of the eigenvalue on the *desired time per order* stock—something that was expected given the immediacy of those parameters to the erosion dynamics. According to the elasticity table, if the propensity for employees to cut corners under work pressure were to be increased, i.e. make *alpha* more negative, this would, as expected, *increase* the weight of the long-term decay behavior mode on *desired time per order* and it would also accelerate the decay rate (note the *positive* elasticity of the eigenvalue to *alpha* in the middle column). The effects of the time constant to adjust desired time per order have similar magnitudes but opposite signs. Increasing the *time to adjust DTO down* would reduce the projection of the behavior mode on *desired time per order* and would reduce the long-term decay rate.

Although in a different order than the original table in the 2016 paper, the updated table contains all the parameters in the first three policies explored by Oliva and Sterman (2001). While they performed their testing using the equilibrium initialization *service backlog* and *desired labor*, Table 1 provides explanations of why policies 1 and 3 would not have been successful even if eliminating the effect of parameter changes on the initial values of stocks. For policy 1, the effect of *time to adjust labor* is very small compared to the other parameters (ranked 8), and they focused on adjusting the *hiring delay*, which has almost no impact on the relevant weight, and ignored *time to adjust desired labor*, which has an effect two orders of magnitude larger. As for policy 3, increasing *rookie effectiveness* does have the desired effect of reducing the rate of erosion of service quality (note the negative sign of the eigenvalue elasticity), but most of this effect is negated by the positive elasticity of the weight to the same parameter.

Note that none of these changes affects the other results presented in Oliva (2016), nor its original message that the tools behind SDA are an efficacious, efficient and effective way to understand and develop policies for realistic system dynamics models.

Conclusion

While it was not yet possible to identify or develop an aggregated metric for the full effect of a parameter in a stock's weight *and* eigenvalue, we believe that the insights gained by the framework presented in Figure 3 will improve

DOI: 10.1002/sdr

the interpretability of the results generated by the SDA toolset. Furthermore, the framework makes clear that assessing a parameter impact on a stock's behavior is a time-sensitive issue as the elasticity of eigenvalues captures the response of the dynamics of the system while the elasticity of weights reflects the instantaneous change in magnitude of the projection.

We have updated the SDA toolset (Naumov and Oliva, 2018) to consistently apply the analysis under the new set of assumptions (even if stocks are initialized from parameters) and present the results of the analysis in a way that facilitates their interpretation under the new framework. The toolset also automates many of the processes that used to require multiple pieces of software and it expedites the exporting of graphs and tables for offline analysis or publication. Finally, for consistency, the SDA toolset's website now includes the updated reports of the analysis of previously published models that were not performed under the assumptions listed here.

Although there is still much work to be done to reduce the cost of performing and interpreting EEA (Kampmann and Oliva, 2017), the updated platform also facilitates the deployment of new ways to present and process the results (see Naumov and Oliva, 2018, for list of planned improvements). We hope that these improvements encourage more modelers to use the toolset and expand the usage and interpretation of EEA. It is only through the accumulated experience of assessing numerous models through this lens that we can further develop SDA's potential.

Biographies

Sergey Naumov is a PhD Candidate in System Dynamics group at MIT Sloan School of Management. He holds an MSc in Engineering and Management from System Design and Management program at MIT and an MSc in Mechanical Engineering from Bauman Moscow State Technical University. His research focuses on sustainable and behavioral operations management with the empirical emphasis on transportation policies and the impact of autonomous and alternative fuel vehicles on mobility services, automotive industry, and the environment.

Rogelio Oliva is the Bob and Kelly Jordan Professor of Business in the Department of Information and Operations Management at Mays Business School at Texas A&M University. He holds a BS in Industrial and Systems Engineering from Monterrey Tech (Mexico), an MA in Systems in Management from Lancaster University (UK), and a PhD in Operations Management and System Dynamics from the Massachusetts Institute of Technology. His research explores how behavioral and social aspects of an organization interact with its technical components to determine the firm's operational performance.

DOI: 10.1002/sdr

References

Duggan J, Oliva R. 2013. Methods for identifying structural dominance. *System Dynamics Review* **29**(Virtual Special Issue). Available: http://onlinelibrary.wiley.com/journal/10.1002/(ISSN)1099-1727/homepage/VirtualIssuesPage.html.

Kampmann CE. 2012. Feedback loop gains and system behavior (1996). *System Dynamics Review* **28**(4): 370–395.

Kampmann CE, Oliva R. 2006. Loop eigenvalue elasticity analysis: three case studies. *System Dynamics Review* **22**(2): 141–162.

Kampmann CE, Oliva R. 2008. Structural dominance analysis and theory building in system dynamics. *Systems Research and Behavioral Science* **25**(4): 505–519.

Kampmann CE, Oliva R. 2017. Analytical methods for structural dominance analysis in system dynamics. In *Encyclopedia of complexity and systems science: living reference work entry*, Meyers R (ed). Springer: New York. https://doi.org/10.1007/978-3-642-27737-5_535-2.

Moxnes E, Davidsen PI. 2016. Intuitive understanding of steady-state and transient behaviors. *System Dynamics Review* **32**(2): 130–155.

Naumov S, Oliva R. 2018. *Structural dominance analysis toolset.* Available: http://iops.tamu.edu/faculty/roliva/research/sd/sda/ [24 July 2018].

Oliva R. 2015. Linking structure to behavior using eigenvalue elasticity analysis. In *Analytical Methods for dynamics modelers*, Rahmandad H, Oliva R, Osgood ND (eds). MIT Press: Cambridge, MA; 207–239.

Oliva R. 2016. Structural dominance analysis of large and stochastic models. *System Dynamics Review* **32**(1): 26–51.

Oliva R, Sterman JD. 2001. Cutting corners and working overtime: Quality erosion in the service industry. *Management Science* **47**(7): 894–914.

Saleh M, Oliva R, Kampmann CE, Davidsen PI. 2010. A comprehensive analytical approach for policy analysis of system dynamics models. *European Journal of Operational Research* **203**(3): 673–683.

Sterman JD. 2000. *Business dynamics: Systems thinking and modeling for a complex world*. Irwin/McGraw Hill: Boston, MA.

DOI: 10.1002/sdr

Appendix: Model structure

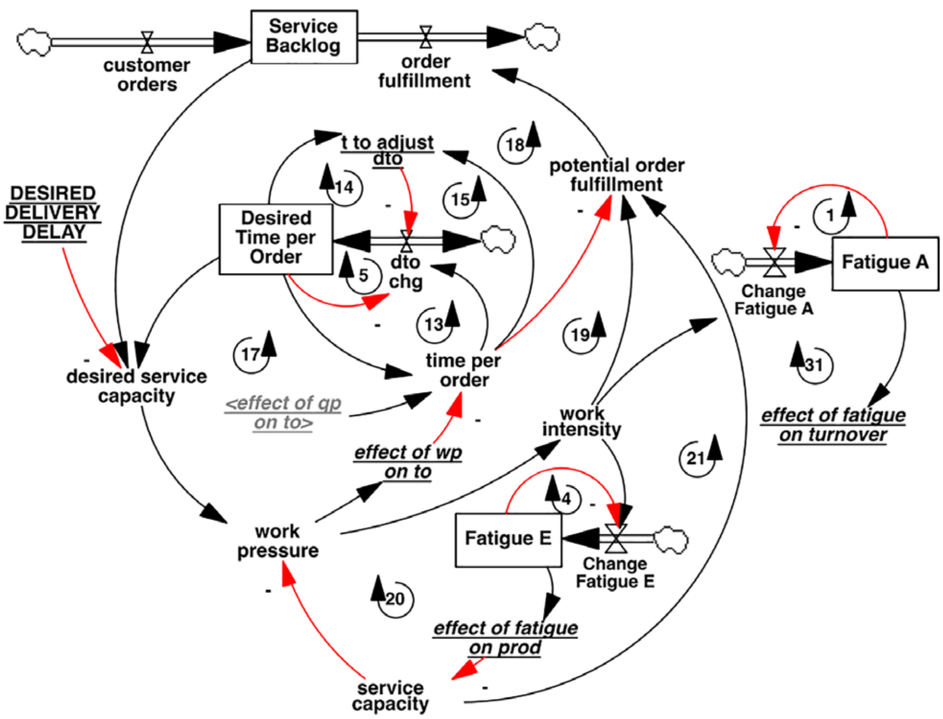

Source: Oliva (2016)

DOI: 10.1002/sdr