In Proceedings of ICCM - 2007- Eighth International Conference on Cognitive Modeling. 61 - 66. Oxford, UK: Taylor & Francis/Psychology Press.

Learning to Control a Dynamic Task: A System Dynamics Cognitive Model of the Slope Effect

Cleotilde Gonzalez (coty@cmu.edu) and Varun Dutt (varundutt@cmu.edu)
Dynamic Decision Making Laboratory
Social and Decision Sciences Department
Carnegie Mellon University, 5000 Forbes Avenue
Pittsburgh, PA 15213 USA

Abstract

We developed a system dynamics model for a simple, but important stock and flows task where the objective was to control the water level in a tank within an acceptable range of the goal, over a number of time periods, in the presence of an unknown environmental inflow and outflow. We also report how this model accounts for human behavior, using behavioral data we collected from human subjects in the task. This exercise helped us understand the strategy and mechanisms our participants used in the simple stock and flows task and develop a model on the task. The model provides an integrated explanation on how the variation in the parameters of the model affects the performance and learning for the participant's task. Finally, we present the model's validity and predictions derived by looking into how the human data fits different learning conditions.

Keywords: Learning; Dynamic decision making; System dynamics model; Stock; Flows.

Introduction

An understanding of the building blocks of dynamic systems, such as stocks (accumulation), flows (rates of incomes and outcomes) and feedback (cause-effect relationships), is essential for dealing with realistic dynamic problems. For example, firms must manage their cash flows to maintain adequate stocks of working capital, and production must be adjusted as sales vary to maintain sufficient inventory. Other examples of these dynamic decision making problems include global warming (Sterman & Sweeney, 2002), factory production, demands and prices of goods (Forrester, 1961), and extinction of natural resources (Moxnes, 2003).

At the individual level, each of us also faces similar stock management challenges: we manage our bank accounts (stock of funds) to maintain a reasonable balance as our incomes (inflows) and expenses (outflows) vary, and we struggle to maintain a healthy weight by managing the inflow and outflow of calories through diet and exercise.

Accumulation is a pervasive process in everyday life, and arises at every temporal, spatial and organizational scale (Cronin, Gonzalez, & Sterman, under review). All stock-flow problems share the same underlying structure: a resource level (stock) accumulates its inflows less its outflows over time.

Unfortunately, there is strong and increasing evidence of poor human understanding of these basic concepts of dynamic systems. For example, Sweeney and Sterman (2000) presented MIT graduate students with a paper problem concerning accumulation of water in a bathtub and asked them to sketch the path for the quantity of water in the bathtub over time, given the patterns for inflow and outflow of water. Despite the apparent simplicity of this task (due to the presence of linearity in the inflow and outflow), they found that only 36% of the students answered correctly. More recently, researchers have found that this misunderstanding of the concepts of flow and accumulation is more fundamental, a phenomenon that has been termed *stock-flow failure* (Cronin & Gonzalez, 2007; Cronin et al., under review). Poor performance in the interpretation of very simple stock and flow problems cannot be attributed to an inability to interpret graphs, contextual knowledge, motivation, or cognitive capacity (Cronin et al., under review). Rather, stock-flow failure is a robust phenomenon that appears to be difficult to overcome.

In past research the stock-flow failure has been investigated through the perception and judgment of static graphs representing flows (Cronin et al., under review; Sterman, 2000). Guided by the tradition of research in dynamic decision making (DDM), we believe that an understanding of the causes and cure for the stock-flow failure will arise through the research on human learning, where individuals can actually experience the flows, influence the stock with their decisions, and have an extended opportunity to control the dynamic system.

With this goal in mind, we constructed a tool, "Dynamic Stock and Flows" (DSF) (Gonzalez & Dutt, submitted), that represents the simplest possible dynamic system containing its most essential elements: a single *stock* that represents accumulation (i.e., water) over time; *inflows*, which increase the level of the stock; and *outflows*, which decrease the level of the stock. We have conducted several human experiments using DSF, including the investigation of environmental functions, the effects of feedback delays, and the timing of decisions, among others.

In this paper we present the results from an initial empirical study aimed at determining the effects of the slope of an inflow on dynamic control of DSF (Dutt & Gonzalez, 2007). Then, we present a system dynamics model that we created to study the cognitive processes involved in dynamic decision making with DSF. We validate the model's results against human data, and present some interesting predictions that emerged from this model. The implications and use of system dynamics modeling of cognitive phenomena are discussed.

The Dynamic Stock and Flows Task

DSF is a generic dynamic control task that we designed to help understand human dynamic decision making, and more concretely for this paper to understand the stock-flow failure.

The full capabilities of DSF are described elsewhere (Gonzalez & Dutt, submitted), and .here we will only give a brief overview of the DSF capabilities relevant for the empirical data and modeling reported in this paper.

The goal in DSF is to reach and maintain the level of water in a tank at a target level over a number of time periods. The level of water in the tank is the stock that increases with the inflows and decreases with the outflows. There are two types of inflows and outflows in this task: those that are exogenous (outside of the decision maker's control) and those endogenous (under the decision maker's control). The exogenous flows are called Environmental Inflow (that increases the level of the stock without the user's control) and the Environmental Outflow (that decreases the level of stock without user's control). The endogenous flows are User's Inflow and Outflow. These amounts are the main decisions made by the user in each time period that increase (user inflow) or decrease (user outflow) the level of the stock.

Figure 1 presents the graphical user interface of DSF. At each time period users see the values of Environment Inflow and Outflow, values of User Inflow and Outflow, the amount of water in the tank (stock) and the goal level. At each time period, users can submit two values (including zero): User Inflow and User Outflow, and click in the submit button. Users may also receive a 'bonus' performance monetary incentive in each time period in which they were close enough to the target level.

Figure 1: A screenshot of the Dynamic Stock and Flows (DSF) task.

The Slope Effect

In dynamic decision making, it has been observed that people may detect linear, positive correlations given enough trials with outcome feedback. However, people have difficulty when there is random error or non-linearity and negative correlations (Brehmer, 1980). As part of the stock-flow failure studies we have also observed that people have difficulty understanding the effects that increasing or decreasing trends of inflow and outflow have on controlling a stock (Gonzalez and Vanyukov, in preparation).

In a laboratory study we investigated how individuals controlled DSF over 100 time periods of practice when the environmental inflow increased (positive slope) or decreased (negative slope) as a function of time period (Dutt & Gonzalez, submitted).

Participants played DSF for 100 time periods with the objective of maintaining the tank's water level at the 4 gallon goal line (within +/- 0.1 gallons). In experiment 1, we used an Environment Inflow function that was either increasing: 0.08 * (*timeperiod*) + 2 or decreasing: (-7.92/99) * (*timeperiod-1*) + 10. Environment Outflow was constant and set at 0 gallons/time period during all 100 time periods. Both the increasing and decreasing functions resulted in an equal amount of environmental net flow into the tank over the course of 100 time periods (604 gallons).

Results showed that the stock was higher for the decreasing function condition (M = 5.909, SE = .205) than the increasing function condition (M = 4.297, SE = .027) ($F(1,31)=12.71$, $p<.001$). The analyses also indicated that the participants' inflow, outflow and stock diminished significantly over time in both conditions (i.e., subjects learned to control the system) ($F(1,31)=9.894$, $p<.001$); and the decrease of the participants' outflow and stock interacted with the slope of the Environmental Inflow function ($F(1,31)=7.031$, $p<.001$) (Figure 2 illustrates the interaction on the stock measure).

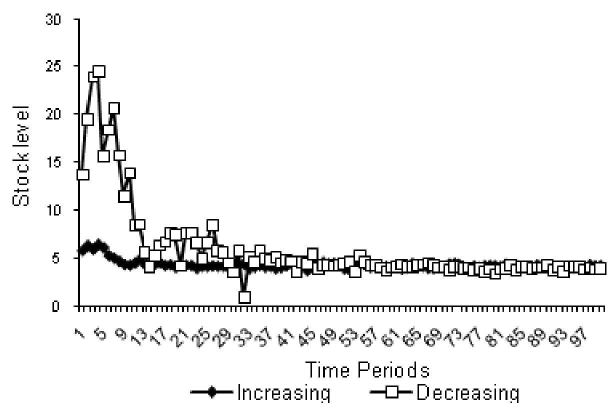

Figure 2: The stock for increasing and decreasing linear Environment Inflow curve conditions over 100 time periods.

In experiment 2, we used a non-linear environment inflow function that was again either increasing: 5*LOG (timeperiod) or decreasing: 5*LOG (101- timeperiod). Outflow was constant and set at 0 gallons/time period during all 100 time periods. Both the increasing and decreasing functions resulted in an equal amount of environmental net flow into the tank over the course of 100 time periods (831 gallons).

Results again showed that the user outflows follow the Environment Inflow functions quite closely. The user inflow and stock indicate that most variability occurred during the first half of the trials, where the decreasing (negative slope) function result in higher inflow and higher stock levels than the increasing (positive slope) function. No difference between the increasing and decreasing functions is observed for the last 50 trials of the experiment in the user inflow and stock results. An analysis of the first 50 trials indicated that the stock was on average higher in the decreasing (M = 7.938, SE = .419) than in the increasing condition (M = 4.757, SE = .143) ($F(1,30) = 6.49$, $p<.05$). The same analyses for the last 50 trials of the experiment did not show a difference between the increasing and decreasing functions for inflow and stock variables. A significant difference was found only for the user outflow (F(1,31) = 85.334, p<.001) where the outflow was higher for the increasing (M = 9.543, SE = .081) than the decreasing (M =7.067, SE = .195) function. Once again, the decrease of the user outflow and stock interacted with the slope of the environmental inflow function; (Figure 3 illustrates the interaction on the stock measure).

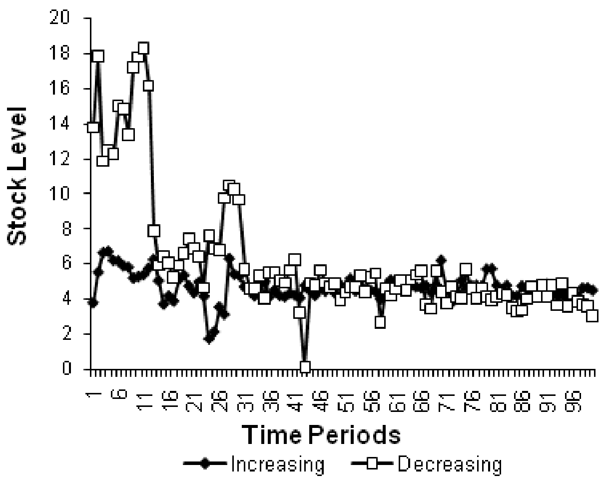

Figure 3: The nature of stock for increasing and decreasing Environment Inflow in non-linear curve slopes over 100 time periods.

System Dynamics Model

System dynamics (SD) is a field that was created by Jay W. Forrester at MIT in the late 1950s and it involves a modeling approach using computer simulations (see Forrester, 1990 for an historical view of the field; Lane, 2000 for a discussion on the modeling approach).

A model in SD involves at its essence the concept of a feedback loop: the collection of information about the system state followed by an action that changes the state of the system (Lane, 2000). These causal links involve delays and non-linearities as well as processes of accumulation (stocks) and flows.

SD modeling has largely focused on the representation of social systems and their evolution over time. In fact, it has been argued that SD is concerned with aggregate social phenomena, not with individual meaningful actions (Lane, 2000).. In this paper we use SD modeling to represent and reproduce the dynamics of individual human behavior found in DSF. We also construct a cognitive interpretation of the SD model, something uncommon in the SD field.

A SD model was developed using Vensim®, an open source modeling software by Ventana Systems Inc. The software has a flexible GUI that provides easy capability to the modeler to represent stocks, define the stocks' inflows and outflows and define their causal relationships. Although the conventions for representing stocks and flows followed in Vensim® are well known and documented in the SD literature (Forrester, 1961; Sterman, 2000) we discuss only some of the many software features that we used for our model of learning and the slope effect. These features are fixed time delay and smoothing.

The function defined in Vensim® as DELAY FIXED (X, T, I) creates a delay of T time periods in an input X with the initial value I of the variable used on the left hand side of the function. In our model we use DELAY FIXED to create a unit time delay in the environment inflow at each time period. This is because participants in our DSF task are aware of the environment inflow value for a time period only at the end of that time period.

Smoothing is defined in Vensim® as SMOOTH (X, T) and SMOOTHI (X, T, I) and creates an exponential smoothing of T time periods in an input X with I as the initial value of the variable used on the left hand side of the function. If X is a step function which jumps to a new value X' at a time instance *t*, then the SMOOTH of X will start from X and approach the value X' over a long range of time periods. The greater the value of T the more time SMOOTH of X takes to approach X'. This smoothing effect of time averages to represent expectations is similar to blending parameters used in learning models of dynamic decision making under the ACT-R cognitive modeling approach (Gonzalez, Lerch, & Lebiere, 2003). In our SD model, we used a smoothing effect to account for the gradual correction of a discrepancy made by participants.

A System Dynamics Model of the Slope Effect in DSF

To help develop this model, we used our observations from verbal protocols collected from four participants (Dutt & Gonzalez, 2007). We also used human data analyses of inflow and outflow decisions and their resulting stock; the averages of individuals' decisions for each of the conditions; and comparisons of the participants' inflow and outflow decisions to the stock and environmental flow values.

Based on these observations and empirical data analyses, we developed the SD model shown in Figure 4. The system essentially consists of 2 inputs (User Inflow and Environmental Inflow) that increase the stock and 2 outputs (User Outflow and Environmental Outflow) that decrease the stock.

The behavior is represented by causal loops described in the model. The Environmental Inflow and Outflow are perceived by the participants. The perception may be different from the reality, as determined by the Environment perception (EP) parameter. Then, the perceived environment netflow is used to forecast the future flow under a Forecast Horizon (FH). This forecast together with the perceived current discrepancy between the stock and the goal are used to determine the netflow correction. This correction is to account for the increase in discrepancy over the perceived time for correction (PTC) smoothed over the memory of discrepancy (MD). Then, according to the determined user netflow correction the user inflow and outflow are entered by the user. Over the course of practice, Users modify the weight they put to inflows and outflows and in general, the empirical data demonstrated that individuals end up realizing that they only need to enter the User Outflow to control for the Environmental Inflow (W=1).

The user net flow correction variable in our model (Figure 4) serves as the main decision function for user inflows and outflows and consists of two parts, discrepancy and forecast of flow. The user net flow correction is given by the procedure: If the discrepancy is beyond an acceptable threshold (.1 above or below the goal), then, attempt to correct for such discrepancy little by little, by smoothing the discrepancy over the perceived time for correction (PTC), and according to the memory of discrepancy (MD). Then, add to this smoothed discrepancy value, the value of the forecasted flow.

Discrepancy is the difference between the goal and stock. The forecast of flow is defined as per the formula: SMOOTH (Perceived Environment Net flow, "Forecast Horizon (FH)") where the Perceived Environment Net flow is a fixed unit delay function of environment inflows and outflows (Vensim® formula: DELAY FIXED ("Environment's Perception (EP)"*(Env Inflow – Env Ouflow), 1, 0).

Hence, the discrepancy over PTC is smoothed by MD and only adds to user net flow correction if the discrepancy is above or below the acceptable range of stock in the tank, else its affect to user net flow correction is zero. The perceived environment net flow is smoothed by parameter FH to makeup the forecast of flow. The environment inflow and outflow act on the stock each time period, with a subject becoming aware of actual values only at the end of that time period. This requirement is realized in our model by using a fixed unit delay as shown in the formula of perceived environment net flow. Participants' accuracy to perceive the environment's inflow and outflow affect is represented by the multiplier EP also present in the delay formula.

As experiment 1 consisted of a linear environment function and experiment 2 consisted of a non-linear environment function we simulated our Vensim® model twice over a course of 100 time periods, once for getting model data for fit to experiment 1's subject data and a second time for getting model data for fit to experiment 2's subject data. During the course of each of the two simulation runs, we varied the values of FH, PTC, MD and W with an increment of 0.05 each. Although each of these parameters could take infinitely many values, their chosen values were inspired from results on experiments 1 and 2 and the availability of human data from both experiments.

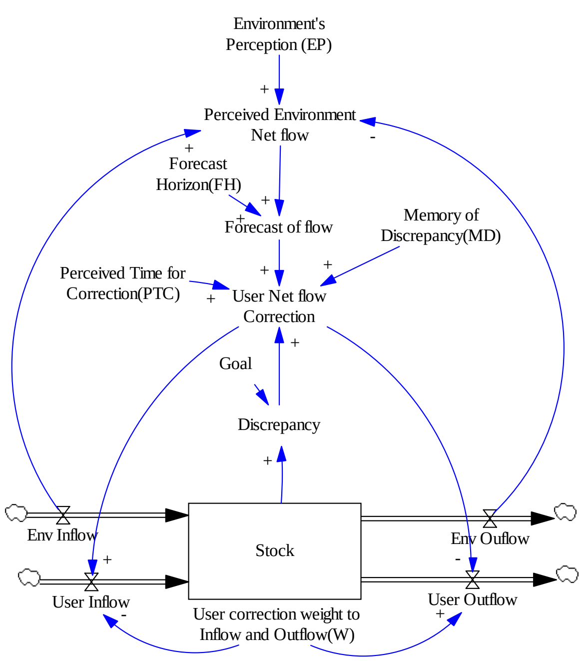

Figure 4: Stock and Flows Model of DSF task.

Fit of Model and Human Data

Based on the results from both experiments, we expected that the parameters of the model, the environment perception (EP), the forecast horizon (FH), the perceived time for correction (PTC), the memory for discrepancy (MD), and the correction for inflow-outflow (W), would be different for the positive and the negative slope conditions.

We went through a process of parameter tuning and sensitivity analysis and were guided by our results and expectations, while measuring the fit of the model's data to the humans' mean and median values over the course of 100 trials, using the r and RMSD statistics (Schunn & Wallach, 2001). The resulting parameter values and data fit statistics for our data from experiments 1 and 2 are summarized in Table 1. The same model depicted in Figure 4 was used to fit the data of the four different data groups from experiments 1 and 2. The value of the parameters summarized in Table 1 help provide a coherent explanation for the different human behavior found between the positive and negative slopes.

Table 1: The value of model parameters and the resulting fit to human data (measured against both, the median and the mean) for each of the 4 groups in the 2 experiments (Linear positive and negative; Non-linear positive and negative) for the stock as dependent variable.

Condition	*Parameters*					*Fit Measures*			
	EP	**FH**	**MD**	**PTC**	**W**	**Median**		**Mean**	
						r	**RMSSD**	**r**	**RMSSD**
				Stock					
Linear Positive L+	1.03	1	1.1	2.5	1	.82	1.94	.79	2.34
Linear Negative L-	1.0	1.5	1.0	4.3	1	.91	3.9	.88	4.15
Non Linear Positive NL+	1.03	1	1	2.0	1	.95	1.85	.51	2.62
Non Linear Negative NL-	1.07	1.5	1	4.5	1	.91	7.54	.63	8.00

Some general observations from results shown in Table 1 are: PTC and FH parameter values are higher in the negative than in the positive functions; the model fits the linear functions better than the non-linear functions; and the model data fit the median of the human data better than the mean.

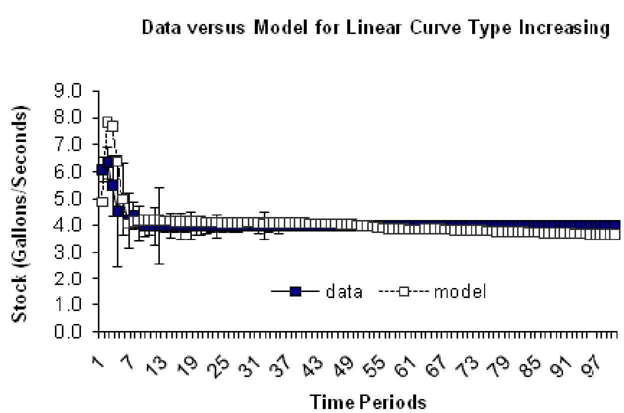

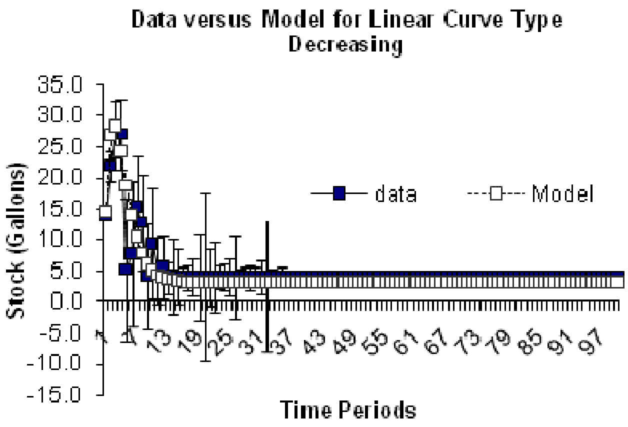

Figure 5: Graphs and parameter values for model's fit to stock.

Figure 5 is a graphical example of the fitting of stock data for the linear and non-linear, positive and negative slopes' functions as a result of our analysis and parameter tuning.

Greater FH value means that participants takes more time to forecast the environment inflow value, and this extra time causes the stock to increase due to the environment inflow, which drives the stock away in each time period. Similarly, greater PTC value means that participants take more time to perceive the discrepancy happening in the tank and hence also take more time to correct the discrepancy; this extra time causes the stock to rise again due to environment inflow into the stock with each elapsing time period. To meet the higher stocks, participants order higher user inflows and outflows, causing user inflow and user outflow values to rise as well. When we fit our model to the negative-sloped environment inflow cases, we see that increasing FH and PTC generates a good fit for the human data. This increase in FH and PTC in our model causes the stock to rise higher (due to the environment inflow action and slower corrective action), causing the discrepancy to increase further and hence, the user net flow correction to increase further. The user net flow correction is further increased by the forecast for net flow, which now happens over a larger time period range (due to FH increase, where there are increased environment inflows over this larger time range). The increase in user net flow correction as described above causes higher user inflow and user outflow in our models, where the user inflow and user outflow are derived from the user net flow correction and weighted by the user correction weight to inflows and outflows (W). The fact that the model's data fits the median of human data more closely than the mean is a reflection of the deterministic nature of SD models in general that fail to account for the variability of human performance, as they are deterministic models.

Model Predictions

From the results on linear and non-linear environmental curves, we find that a negative slope in the environment inflow produces higher stock, user inflow and user outflow as well as higher variability in stock, user inflow and user outflow, particularly in the first 50 time periods (this is also the time when subjects are learning to control the stock in DSF).

If the parameters that we have proposed in our model and their fit to human data can have a cognitive interpretation as we presented in the previous section, then there are some interesting questions we could answer by looking into model fitting more deeply.

Specifically, we are interested in determining how human data from each of the collected groups: Linear Positive (L+), Linear Negative (L-), Non-linear Positive (NL +), and Non-linear Negative (NL -) would fit to the model's data from the other groups. For example, we test how human data in L+ fits to the model's data from L-, NL+, and NL- groups by looking at the difference between the r values (measures of fit) between equivalent groups from the model's data and the comparison group.

For this fit exercise, we found the mean and median of our human data for both experiments 1 and 2 under different conditions, L+, L-, NL+ and NL-. We already have our fits and model parameters calculated as a result of parameter tuning and sensitivity analysis to the mean and median of human data under L+, L-, NL+ and NL- conditions as mentioned in Table 1. We took the model parameters under the NL- condition and fit the model data due to these parameters to the mean of human data under the L- condition. Similarly we took the model parameters under L- condition and fit the model data due to these parameters to the mean of human data under the NL- condition. We did a similar exercise for L+ to NL+ and NL+ to L+. This process helps us to foresee the mapping of our model's cognitive parameters to different experimental conditions as measured by model fits, i.e. how different values of cognitive parameters EP, FH, PTC, MD and W perform under different experimental conditions. This mapping exercise can help us predict how the model experiencing a linear inflow would behave when put into a non-linear inflow and vice-versa. For example, from a managerial perspective a firm may suddenly face non-linear changes in demand after operating under a constant (linear) demand (Paich & Sterman, 1993). In addition, the knowledge gained from this cross fitting exercise helps us understand the nature of underlying task situations involved, task complexity as it would possibly be experienced by the decision makers.

The results from fitting model's parameters on linear environment inflow to non-linear environment inflow and vice-versa are tabulated in Table 2. The *r* values given in Tables 1 and 2 provide results for the mean stock in DSF.

The results show that *r*(NL- to L-) > *r*(L-), *r*(L- to NL-) < *r*(NL-) and *r*(NL- to L-) > *r*(L- to NL-). Also from Tables 1 and 2, similar results hold for the positive slope environment inflow cases for both the linear and non-linear curve types. This means *r*(NL+ to L+) > *r*(L+), *r*(L+ to NL+) < NL+ and *r*(NL+ to L+ > L+ to NL+.

Table 2: Values of correlation coefficients for model predictions on linear and non-linear positive and negative sloped environment inflow for stock mean.

Curve Type Environment Inflow Of Model	Curve Slope Environment Inflow Of Model	Condition	Mean Stock (r)
Linear	Negative	Linear Negative Model parameters and fit to non linear negative human data (L-ve to NL-ve)	0.63
Non Linear	Negative	Non Linear Negative Model parameters and fit to linear negative human data (NL-ve to L-ve)	0.91
Linear	Positive	Linear Positive Model parameters and fit to non linear positive human data (L+ve to NL+ve)	0.50
Non Linear	Positive	Non Linear Positive Model parameters and fit to linear positive human data (NL+ve to L+ve)	0.83

These cross fit results indicate the *r* value diminishes from linear to non-linear curves but increases from non-linear to linear inflows. This is, the non-linear conditions are more difficult to fit than the linear condition.

Similar results were reported on the human data collected from both experiments 1 and 2 earlier, where subjects' performance was poorer in the non-linear environment inflow DSF task when compared to performance in the linear environment inflow DSF task. These similarities between the nature of DSF model and the DSF human data also further support the choice of the model's parameters.

Conclusions

Stock-flow failure is a phenomenon representing the poor interpretation of very simple problems involving accumulation over time by flows (Cronin & Gonzalez, 2007). In this paper we investigated one possible explanation for the stock-flow failure and that is the increased difficulty for controlling systems with decreasing more so than increasing inflows. We investigated this simple failure in a Dynamic Stock and Flows (DSF) task. We found that participants yielded greater quantity of stock, inflows and outflows and more variability in them for negative slope (decreasing) environment inflow conditions when compared with the quantity and the variability in stock inflows and outflows for the positive slope (increasing) environment inflow conditions.

We explain these results through a system dynamics model that helps derive differences in human behavior due to slopes of environment inflow. The constructed model and its fit revealed minimize of a number of human cognitive parameters in the dynamic task which makes us think that similar cognitive parameters would constitute many such simple dynamic stocks and flows tasks which are important in our day to day lives (our bank accounts to our weight gain and loss processes to name a few) where the understanding of such parameters would be helpful in overcoming the difficulties that most of us face while encountering them.

Acknowledgments

This research was partially supported by the National Science Foundation (Human and Social Dynamics: Decision, Risk, and Uncertainty, Award number: 0624228) and by the Army Research Laboratory (DAAD19-01-2-0009) awards to Cleotilde Gonzalez.

References

Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica, 45*, 223-241.

Cronin, M., & Gonzalez, C. (in press). Understanding the building blocks of system dynamics. *System Dynamics Review*.

Cronin, M., Gonzalez, C., & Sterman, J. D. (under review). Why don't well-educated adults understand accumulation? A challenge to researchers, educators and citizens.

Dutt, V., & Gonzalez, C. (2007). *Slope of inflow impacts dynamic decision making*.

Forrester, J. W. (1961). *Industrial dynamics*. Waltham, MA: Pegasus Communications.

Forrester, J. W. (1990). *The Beginning of System Dynamics*. Boston, MA: The Sloan School of Management, MIT.

Gonzalez, C., & Dutt, V. (submitted). A generic dynamic control system for management education.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science, 27*(4), 591-635.

Lane, D. C. (2000). Should system dynamics be described as a 'hard' or 'deterministic' systems approach? *Systems Research and Behavioral Science, 17*, 3-22.

Moxnes, E. (2003). Misperceptions of basic dynamics: The case of renewable resource management. *System Dynamics Review, 20*, 139-162.

Paich, M., & Sterman, J. D. (1993). Boom, bust and failures to learn in experimental markets. *Management Science, 39*(12), 1439-1458.

Schunn, C. D., & Wallach, D. (2001). Evaluating goodness-of-fit in comparison of models to data. University of Pittsburgh.

Sterman, J. D. (2000). Learning in and about complex systems. *Reflections: The SoL Journal, 1*(3), 24-51.

Sterman, J. D., & Sweeney, L. B. (2002). Cloudy skies: Assessing public understanding of global warming. *System Dynamics Review, 18*(2), 207-240.

Sweeney, L. B., & Sterman, J. D. (2000). Bathtub dynamics: initial results of a systems thinking inventory. *System Dynamics Review, 16*(4), 249-286.