CHAPTER 15

USING SIMULATION TO ANALYZE SIMPLE POSITIVE AND NEGATIVE LOOPS

The computer simulation techniques developed so far can be used to analyze the behavior of simple positive and negative loops. Because simple positive and negative loops form the building blocks of more complex models, it is important to understand the kinds of behavior they can generate in fairly rich detail. The examples that follow illustrate some of the most common positive and negative loop structures, and some useful ways of employing simulation to probe system behavior.

EXAMPLE I: YEAST BUDDING (POSITIVE LOOPS)

The simplest and most fundamental positive feedback loop consists of one level and one rate, and the rate is directly proportional to the level. An example, shown in Figure 15.1, is the model of yeast budding taken from Chapter 14. (The equations for this model are listed after Exercise 3 of Chapter 14.)

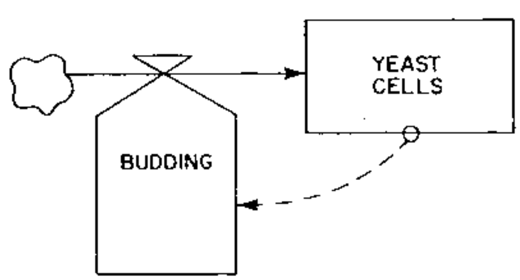

Figure 15.1 Flow diagram of yeast budding

In a simple positive feedback loop with one rate and one level, as the level increases, the rate increases as well, so the level grows at an increasing pace. If the rate is directly proportional to the level (as in the yeast example), the behavior generated is exponential growth.

A quantity that is growing exponentially will double in a fixed amount of time, no matter how long it has been growing or how large the quantity has become. For example, in the yeast budding loop, if the budding fraction BUDFR = 0.1, it can be shown that the doubling time for the number of yeast cells is roughly 7 hours. Thus, if the initial number of yeast cells is 10, the number of cells will reach 20 in about 7 hours, and it will reach 40 in another 7 hours.[1]

The following exercises provide an opportunity to explore the relationship between the growth fraction and the doubling time in simple positive loops.

Exercise 1: Simulation of a Positive Feedback Loop

Look up the equations for yeast growth from Chapter 14. From your first run of the model, describe the behavior of the level and rate.

Exercise 2: Doubling Time

a. Run the yeast model for forty simulated hours.

b. Measure the time needed for the number of yeast cells to double from the initial value of yeast.

c. Measure the time required for the number of yeast to double from its value at hour twenty.

Exercise 3: Effect of Budding Fraction

a. Run the yeast model with a budding fraction of 0.2.

b. Run the model with a budding fraction of 0.05. How do the results differ? Does the model still generate exponential growth? Is there a value of the budding fraction which will cause the model not to produce exponential growth?

c. If the budding fraction were negative, would the feedback loop still be positive?

Exercise 4: The Bank Account—Part I

Suppose you deposit $500 in a bank account earning 10 percent interest compounded annually.

a. Draw a causal-loop diagram and flow diagram for the bank account case. (Assume no money is withdrawn from the account.)

b. Write DYNAMO equations and simulate the bank account for a twenty-year period. (Set DT = 1 year).

c. How much money is in the account in year twenty? What is the doubling time for the account?

Exercise 5: The Bank Account—Part II

In Part I of the bank account problem, you deposited $500 in the bank and left it there to gather interest. Suppose the account earns interest exactly as before, but you must withdraw at the constant rate of $50 per year from the account.

a. Modify your flow diagram to include the withdrawal rate of $50 per year.

b. Modify your DYNAMO equations and run the model on the computer. How do the results differ from your results in Part I?

c. Suppose you begin with $600 in your account, rather than $500, and withdraw $50 per year. How do the results differ?

d. Suppose you begin with $400 in the account. How do the results differ?

STARTING A MODEL IN EQUILIBRIUM

In analyzing the behavior of a system, it is often helpful to begin by determining the system's equilibrium point. This can be done by trial and error, but it is often easier to determine the equilibrium point by examining the flow diagram and system equations. For example, in the preceding bank account case, equilibrium occurs when the money withdrawn each year exactly equals the amount of interest earned. If $50 is withdrawn per year, this means that equilibrium occurs when $50 interest is earned. If the rate is 10 percent, this corresponds to a bank account balance of $500.

Once you have determined the equilibrium point mathematically, it is easy to check your calculations by simulating the results on the computer. Just set the initial values of the system levels to their equilibrium points. If your calculations are correct, the model should remain in equilibrium.

Exercise 6: Calculating Equilibrium Values

a. Suppose you withdraw $60 per year from a savings account earning 10 percent interest. What is the equilibrium balance?

b. Suppose you withdraw $50 per year from an account earning 8 percent interest. What is the equilibrium balance?

EXAMINING A SYSTEM'S RESPONSE TO DISTURBANCES—PART I

Once you have started a system in equilibrium, it is often useful to see how the system responds to exogenous (i.e., outside the system) disturbances. For example, suppose you place $500 in a bank account earning 10 percent interest, and withdraw $50 per year. Then, however, starting five years from now, it becomes necessary to withdraw $75 per year. How will the system respond?

One way to test the response of the system is to use a special DYNAMO function called the STEP function to simulate the sudden $25 increase in the withdrawal rate beginning in five years. A reasonable set of equations for the model includes the following:

```
*    BANK ACCOUNT
L    BAL.K=BAL.J+(DT)(INT.JK-WDRW.JK) DOLLARS
N    BAL=500
R    INT.KL=(0.10)(BAL,K) DOLLARS/YEAR
R    WDRW.KL=50+STEP(25,5) DOLLARS/YEAR
```

The expression "STEP(25,5)" instructs the computer to increase the withdrawal rate by 25 dollars in year 5. Thus a graph of the withdrawal rate would look like Figure 15.2.

The STEP function can be used whenever it is necessary to simulate a sudden step change in a system rate. The general DYNAMO form of the STEP function is:

STEP(HEIGHT,STTIME)

where HEIGHT is the height of the STEP, and STTIME is the abbreviation for step time, the time when the step occurs. Any numerical values can be used for HEIGHT and STTIME.

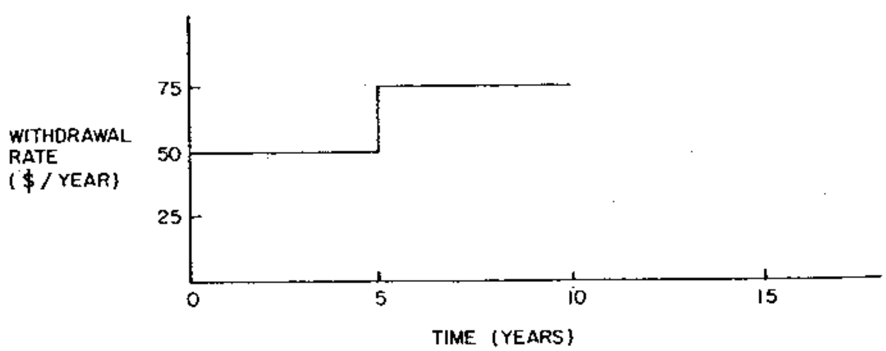

Figure 15.2 Step function graphed

Exercise 7: Using the STEP Function

a. Use the STEP function to test the response of the bank account model to an increase in the withdrawal rate from \$50 to \$75 at TIME = five years. What behavior does the system generate?

b. Use the STEP function to test the response to the model to a decrease in the withdrawal rate from \$50 to \$30 at TIME = three years. What behavior does the system generate?

EXAMPLE II: YEAST DEATHS (NEGATIVE LOOP)

The simplest and most fundamental negative loop contains one rate and one level. An example, shown in Figure 15.3, is the yeast deaths case taken from Chapter 14. If, in a simple negative loop, the rate is directly proportional to the level, the loop will generate exponential decay.

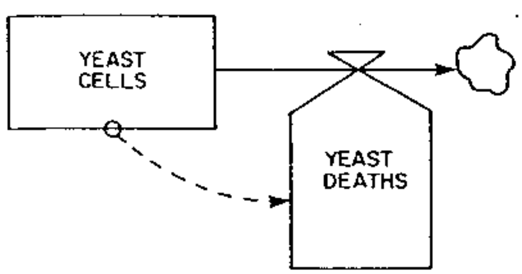

Figure 15.3 Yeast deaths

A quantity that is decaying exponentially will move half-way to its equilibrium value in a fixed amount of time, no matter how far from equilibrium it begins. For example, if the average lifetime of yeast is 20 hours, it can be shown that the halving time is 14 hours. Thus if the initial number of yeast cells is 10, the number of cells will fall to 5 in 14 hours; and it will fall to 2.5 in another 14 hours.[2]

The exercises that follow provide an opportunity to explore the relationship between the average lifetime and the halving time for simple negative loops.

Exercise 8: Simulation of a Negative Loop

Simulate the yeast deaths loop, using the equations developed in Chapter 14. (Assume an initial value of 10 yeast cells, and do not include yeast budding.) Examine the behavior of the level and the rate. How do they differ from the behavior of the yeast budding loop? What is the equilibrium point for the number of yeast cells?

Exercise 9: Halving Time

How long does it take the number of yeast cells to get half-way from its initial value to its equilibrium value? How long does it take it to get from half-way to one-quarter of the way?

Exercise 10: Effect of the Yeast Lifetime

Run the model with an average lifetime of yeast equal to 10 hours. Run the model again with the average lifetime of yeast equal to 40 hours. Does the behavior in each case still represent exponential decay? How do the results differ?

Exercise 11: Yeast Model with Budding and Deaths

Run the yeast model, including both yeast budding and yeast deaths. Does the model exhibit exponential growth or decay? If the model exhibits growth, try to find values of constants that will cause the model to show decay. If the model shows decay, try to find values of constants that will cause the model to show growth. Is there a set of constants that will cause the model to show growth and then decay? Why or why not?

Exercise 12: Central Library—Part I

Books in the Central Library in East Rapids are frequently stolen or lost, and often they just plain fall apart. In fact, the average lifetime of books in the library is just 10 years. The East Rapids City Council provides a library budget large enough for the purchase of 500 new books a year.

- **a.** Draw a causal-loop diagram and flow diagram for the Central Library case.
- **b.** Write equations for the model.
- **c.** Determine the equilibrium point for the number of books in the library, and start the model in equilibrium.

Exercise 13: Central Library—Part II

The City Council in East Rapids has just completed a lengthy analysis of its budget, and, as a result, the library budget is expected to fall sharply, starting in three years. With the planned budget cut, the library will be able to purchase only 300 books a year, rather than 500.

a. Use a STEP function to simulate the effects of the sharp reduction in the book purchase rate in three years.

b. What is the new equilibrium level of books? How long does it take the number of books to fall half-way to the new equilibrium level?

c. Suppose a wealthy donor provides an endowment that permits the library to purchase 700 books a year, beginning in six years. Use a STEP function to simulate the increase in the purchase rate. What is the new equilibrium? How long does it take the number of books to rise half-way to the new equilibrium?

d. How would the number of books in the Central Library be affected, if the average lifetime of books could be increased from 10 years to 20 years?

MORE COMPLEX RATE FORMULATIONS

The simple positive and negative loops considered so far all involve rate equations based on one of the following three forms:

RATE = CONSTANT
RATE = LEVEL*(GROWTH FRACTION)
RATE = LEVEL/(AVERAGE LIFETIME)

The following example introduces a rate formulation that is closely related to the average lifetime formulation, but is more complex.

EXAMPLE III: COFFEE COOLING (NEGATIVE LOOP)

This example involves performing a simple physical experiment and then modeling the system in the experiment. You will need the following equipment:

1. A cup of hot coffee or other liquid at a temperature well above room temperature;
2. A thermometer capable of measuring the temperature of the coffee (use one that will read at least as high as 100 degrees Celsius; a laboratory or candy thermometer should work);
3. A watch or stop watch;
4. A pencil and paper.

While the temperature is still hot, measure the coffee's temperature at regular intervals (every few minutes). Record the temperature and the time of the reading on a sheet of paper. Stop taking readings when the temperature is near room temperature.

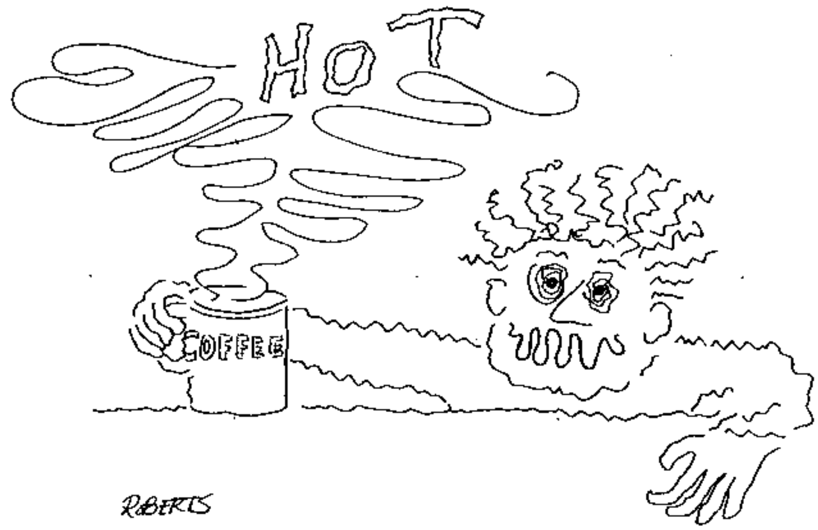

Exercise 14: Graphing Temperature

Plot the temperature versus time.

MODELING TEMPERATURE CHANGE

Development of a model of coffee cooling starts, of course, with a causal-loop diagram. Figure 15.4 depicts a possible diagram.

The diagram says that the decline in temperature reduces the temperature, and the lower the temperature, the smaller the decline. (Is this hypothesis consistent with the data from your experiment?) Figure 15.5 depicts a flow diagram based on the causal-loop diagram in Figure 15.4.

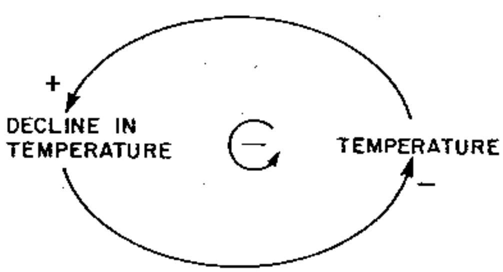

Figure 15.4 Causal-loop diagram of temperature change

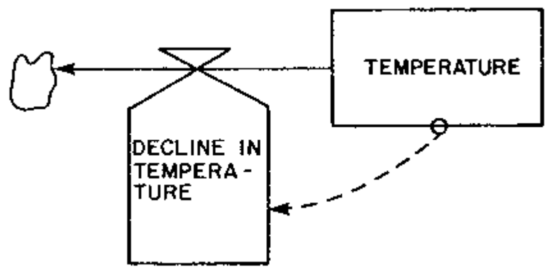

Figure 15.5 Flow diagram of temperature change

Using the name TEMP for temperature (in degrees Celsius) and the name DECLINE for the decline in temperature (in degrees Celsius per minute), the level equation for temperature can be written in the usual form:

TEMP.K = TEMP.J + (DT) (− DECLINE.JK)

Now, how should the rate equation for the decline in temperature DECLINE be formulated? Two assumptions are involved. First, according to the causal-loop diagram, the higher the temperature, the faster the decline. Second, common sense suggests that, when a cup of coffee reaches room temperature, it will not decline further. Therefore, perhaps the decline in temperature should be formulated as a function of the *difference* between the temperature of the coffee and room temperature. The further above room temperature a cup of coffee is, the faster its temperature will decline.

This suggests the following rate equation:

DECLNE.KL = (TEMP.K − ROOMTP)/T

where ROOMTP is room temperature, and T is a "cooling constant."

The "cooling constant" T determines how fast the adjustment of temperature occurs. (Thus T is measured in minutes.) The larger the value of T, the slower the decline in temperature. The value of T might depend on several things. One is the type of coffee container. For example, a glass generally will release heat more quickly than a ceramic cup. (You might redo the previous exercise, comparing several containers.) In addition, the larger the surface area of the container relative to the volume of liquid, the smaller T generally will be.

Exercise 15: Simulating the Coffee Cooling Case

Write DYNAMO equations and simulate the coffee cooling system. (Set DT = 1 minute.) Once you have simulated the values of temperature, plot them on the graph you used for Exercise 14. To carry out the simulation, you need to estimate T. There are at least two ways to do this. One is to start by simulating the system using a guess for T. If the simulated temperature drops faster than in the experiment, increase T and redo the simulation. Similarly, if the simulated temperature drops more slowly than the experimental data, decrease T. Keep trying values of T until the behavior of the model matches the behavior in the experiment.

Another way to estimate T is as follows. Let TEMP.J equal one observation in the experiment, let TEMP.K equal the next observation, and let DT equal the time between observations. Then T can be estimated using the formula:

T = (DT) (TEMP.J − ROOMTP)/(TEMP.J − TEMP.K)

For example, if a reading at one point is 50 degrees Celsius and two minutes later is 45 degrees Celsius, and if room temperature is 20 degrees Celsius, then T would be:

(2)(50 − 20)/(50 − 45) = 12

Averaging two or more estimates will give a more accurate value of T.

AUXILIARY VARIABLES

The coffee flow coding diagram in Figure 15.5 has one defect: much of the detail involved in the rate equation for the decline in coffee temperature is hidden in a singled dashed line connecting the level and the rate. One way to clarify the flow diagram is to define a new variable DIFF, which is the difference between the temperature of the coffee (TEMP) and the room temperature (ROOMTP): DIFF.K = TEMP.K − ROOMTP. This new variable is called an *auxiliary* variable because it aids in forming a rate. (This is similar to the role of auxiliary verbs in English, which aid in expressing an action verb.) Using the auxiliary variable DIFF, the equations for the coffee cooling model can be rewritten as follows:

```
L   TEMP.K = TEMP.J + (DT)(−DECLNE.JK)   DEGREES
N   TEMP = 50
R   DECLNE.KL = DIFF.K/T   DEGREES/MINUTE
C   T =   MINUTES
A   DIFF.K = TEMP.K − ROOMTP   DEGREES
C   ROOMTP = 20   DEGREES
```

Figure 15.6 indicates how DIFF can be added to the coffee cooling flow diagram.

Auxiliary variables are often useful in formulating complex rate equations. Auxiliaries arise when the formulation of a level's influence on a rate involves one or more intermediate calculations—similar to the calculation of DIFF in the coffee cooling case. This sort of intermediate calculation occurs quite frequently in complex models, as will be seen, and the use of auxiliaries

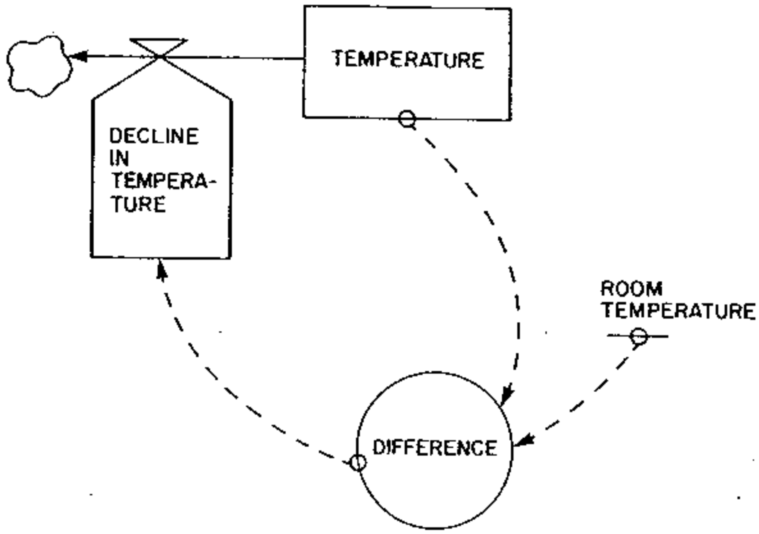

Figure 15.6 Flow diagram with auxiliary added

can clarify otherwise confusing formulations. A summary of flow diagrams symbols is given at the end of this chapter.

THE GOAL-GAP FORMULATION

The rate equation for the decline in temperature in the coffee-cooling model is an example of a quite general rate formulation:

RATE = (LEVEL − GOAL)/(ADJUSTMENT TIME)

In the coffee cooling example, the coffee temperature (the system level) can be viewed as drifting toward a "goal" (the room temperature). The "adjustment time" in this formulation plays a role analogous to the average lifetime in a simple decay formulation: It determines how rapidly the system level adjusts toward its goal. In fact, the simple decay formulation can be viewed as a special case of the goal formulation, in which the "goal" is zero.

RATE = LEVEL/(ADJUSTMENT TIME)

The goal formulation is frequently useful in representing purposeful action. For example, suppose the manager of a department store wishes to maintain a certain fixed number of shoes in stock. The rate at which the manager orders new shoes might depend on the difference between the actual number of shoes in stock and the goal (i.e., the number the manager desires). Furthermore, if a gap exists between the number of shoes desired and the actual number in stock, the manager might not attempt to fill the gap all at once, but might prefer to close the gap gradually, over a period of time.

In this case, the shoe order rate might be written:

ORDERS.KL = (GOAL − SHOES.K)/ADJT

This indicates that the number of shoes ordered per month (ORDERS) depends on the gap between the desired number of shoes in stock (the GOAL) and the current level of shoes in stock (SHOES). Furthermore, the gap is not closed all at once, but instead is closed over a period of time, the adjustment time ADJT. (Note that, in this case, the order rate is formulated "GOAL minus LEVEL." In the coffee cooling example, the decline rate was formulated "LEVEL minus GOAL." The two formulations are equally useful. Which is chosen depends on the logic of the particular example.)

More generally, the goal-gap rate formulation is useful whenever an identifiable *goal* exists, alternatively seen as an objective, a target, a norm, or a desired condition. In comparison or in contrast to this goal is the actual *situation,* which is inevitably a level. The *difference* between the goal and the actual condition, the "goal-gap," is the motivator or driving force underlying corrective action. But the corrective action does not occur all at once; instead, it occurs over some adjustment time. (See Figure 15.7.)

The following exercise provides an opportunity to explore the goal-gap formulation.

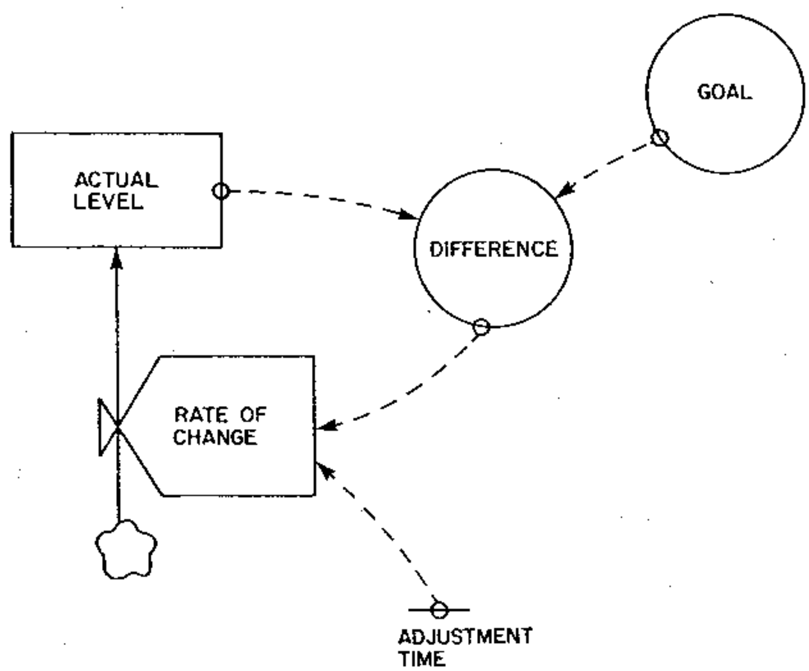

Figure 15.7 Flow diagram for goal-gap formulation

Exercise 16: Jobs and Migration—Part I

According to Example III in Chapter 3, the availability of job openings influences workers to migrate into the city; and as workers migrate into the city, they fill the available openings. This suggests the negative loop shown in Figure 15.8.

a. Formulate a flow diagram for the jobs and migration case, choosing a goal-gap structure. (*Hint:* It is easiest to view the population of workers in the city as the system level.)

b. Write equations for the model. (*Note:* You do not need to choose parameters or run the model on the computer. That will be treated in the last exercise of this chapter.)

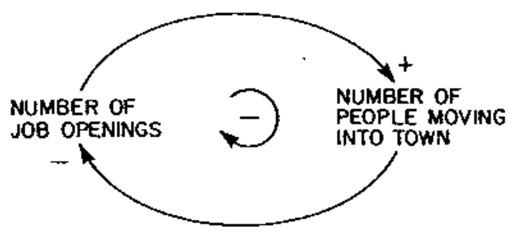

Figure 15.8 Effect of jobs on migration

EXAMINING A SYSTEM'S RESPONSE TO DISTURBANCES—PART II

Recall that earlier in the chapter, the response of the bank account system to an exogenous change in the withdrawal rate was examined, as well as the response of the library system to an exogenous change in the book acquisition rate. In a similar fashion, the response of the coffee cooling system to an exogenous change in room temperature can be analyzed.

For example, suppose a cup of tepid (20 degree Celsius) coffee that has been sitting on the kitchen table is suddenly placed in the refrigerator. How would it respond?

Once again, an easy way to examine the response of the system is to use a STEP function. In this case, the "variable" that needs to be stepped is the room temperature ROOMTP, which is currently a constant. In order to produce a step change in ROOMTP, it is necessary to redefine ROOMTP as an auxiliary variable and give it a subscript K. (ROOMTP must have a subscript, because it now will vary with time.) This produces the following equations:

```
L    TEMP.K = TEMP.J + (DT)(−DECLNE.JK)
N    TEMP = 20
R    DECLNE.KL = DIFF.K/T
C    T =
```

A DIFF.K = TEMP.K − ROOMTP.K
A ROOMTP.K = 20 + STEP(TCHG, 10)
C TCHG = −15

The equation for ROOMTP indicates that after 10 minutes, the room temperature drops 15 degrees, from 20 to 5.

Exercise 17: Stepping Room Temperature

a. Modify the coffee cooling-model to include the STEP change in ROOMTP. Run the model with step change from 20 degrees to 5 degrees as just mentioned.

b. Suppose the coffee was moved from the kitchen table to the sauna (35 degrees centrigrade). How would the system respond?

ADDITIONAL RATE FORMULATIONS

Sometimes, when trying to move from a causal-loop diagram to a flow diagram, it is difficult to decide which variables are levels, which are rates, and which are auxiliaries. In addition, once rates and levels have been identified, it is sometimes difficult to decide on an appropriate formulation for the rate equations. Frequently, none of the formulations discussed so far seem appropriate, and new formulations must be invented to fit the purpose. Often, formulating rate equations requires a certain amount of ingenuity—and a healthy willingness to try out alternative possibilities. The following example illustrates some of the problems involved in writing equations for a somewhat difficult model.

EXAMPLE IV: PUSHUPS AND PRACTICE (POSITIVE LOOP VERSION)

The causal-loop diagram in Figure 15.9 relates the number of pushups Jim can do, and the amount he practices. (See Exercise 1 in Chapter 4 for a discussion of this example.)

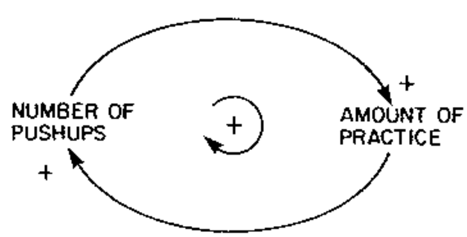

Figure 15.9 Pushups and practice cycle

Both the number of pushups Jim can do and the amount he practices seem at first glance to be levels; certainly neither appears to be a rate. This creates a problem, though, because levels cannot change unless there are rates to change them. Thus a rate must be "hidden" in the causal-loop diagram. One solution to the problem is the following. Assume that the number of pushups Jim can do is a level. Then there ought to be an "improvement rate" that causes the number of pushups to increase. (Jim's improvement rate would be the number of additional pushups he can do per month.) This produces the causal-loop diagram shown in Figure 15.10.

Now, what about "Amount of Practice"? One simplifying assumption would be that the amount Jim practices is a direct function of the number of pushups he can do; and the more he practices, the faster he improves. Under this simplification, "Amount of Practice" would be an auxiliary variable, and the complete flow diagram would look like Figure 15.11. Of course, more complicated assumptions about decisions affecting practicing would lead to quite different flow diagrams.

Now all that remains is to write the equations. According to the flow diagram, the number of pushups Jim can do influences the amount he practices. But what is the exact relationship between the two? Jim would have to be observed for some period to find out. Since this is a speculative model, a simple plausible relationship will be hypothesized. (In fact, one reason to build a simulation model is to examine the implications of plausible hypothesized relationships.)

One simple assumption is that the amount Jim practices is a linear function of the number of pushups he can do. For example, perhaps he practices one-half minute per day for each pushup he can do. This produces the following equation:

Amount of Practice(minutes) = 0.5(minutes/pushup)*Number of Pushups

Thus, for example, if Jim can do 30 pushups, he practices 15 minutes a day, under this assumption.

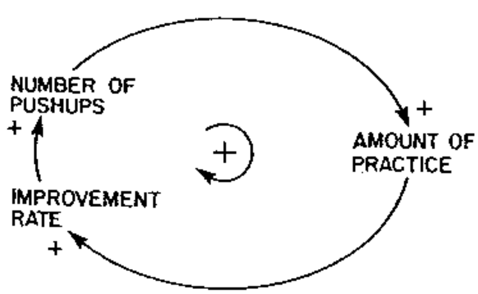

Figure 15.10 Pushups and practice—enlarged

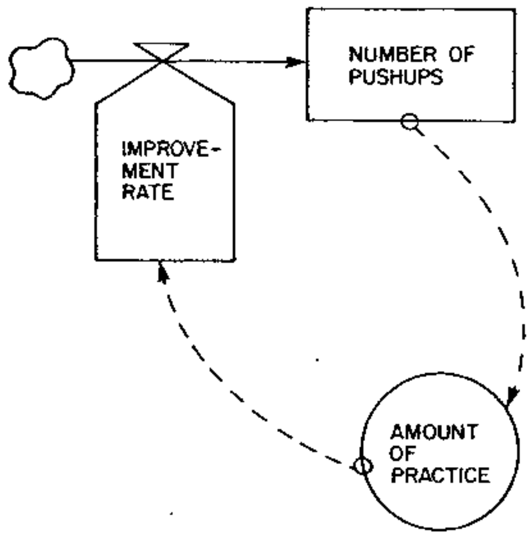

Figure 15.11 Flow diagram of pushups and practice

Now, the relationship between the amount Jim practices and his improvement rate must be formulated. It seems plausible to assume that Jim must practice at least a certain amount of time, simply to maintain his current level of pushups. Any practice above and beyond this maintenance amount would result in improved performance. It might be reasonable to assume that Jim must practice 10 minutes a day to maintain his performance; and for every minute he practices above 10, he improves at the rate of 0.2 pushups per month. This produces the following equation:

Improvement Rate (pushups/month) = (Amount of Practice − 10 minutes)*0.2 (pushups/month/minute)

This relationship implies that when Jim practices 10 minutes a day, he does not improve at all. If he practices 15 minutes a day, he improves at the rate of 1 pushup per month; and if he practices only 5 minutes a day, his performance drops by 1 pushup per month.

Exercise 18: A Model of Pushups and Practice

a. Write DYNAMO equations for the pushup and practice model. (Choose DT = 1 month.)

b. Run the model setting the initial number of pushups Jim can do to 30. Rerun the model, setting the initial number he can do to 20. Rerun the model again, setting the initial number of pushups to 10. How do the results differ? Why?

c. How does the model's equilibrium point depend on the amount of time Jim practices per pushup he can do? Run the model several times, choosing different values for this parameter, "Practice Time per Pushup."

d. How does the equilibrium point depend on the value 0.2 in the equation relating the amount Jim practices and his improvement rate? Try running the model with alternate values for this "Practice Effectiveness" parameter.

e. How does the equilibrium point depend on the value 10 in the equation relating the amount Jim practices and his improvement rate? Run the model with alternate values of the "Maintenance" parameter.

EXAMPLE V: PUSHUPS AND PRACTICE (NEGATIVE LOOP VERSION)

An alternative model can be formulated by assuming that Jim has a pushup goal. Assume that Jim would like to be able to do 50 pushups, and assume as well that the amount he practices is a direct function of how far he is from his goal. As before, assume that Jim's improvement rate depends on the amount he practices. This produces the causal-loop diagram shown in Figure 15.12.

Notice that this goal-gap formulation produces a negative feedback loop, which controls the amount of practice in an effort to achieve a goal of 50 pushups. In Example IV, on the other hand, the loop was positive, potentially resulting in unlimited growth in Jim's ability to do pushups.

The new flow diagram looks like Figure 15.13. (As before, still more complicated or different assumptions about decisions affecting practice would lead to different flow diagrams.)

Now all that remains is to write equations. According to the flow diagram, the difference between the number of pushups Jim would like to do (his goal) and the number he actually can do (the "actual") influences or activates the amount he practices. Jim would have to be observed for some period of

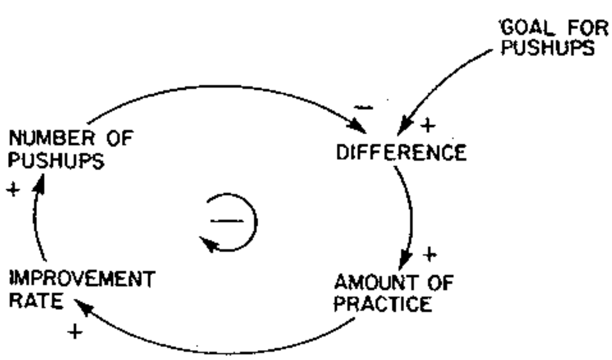

Figure 15.12 Pushups and practice, with goal

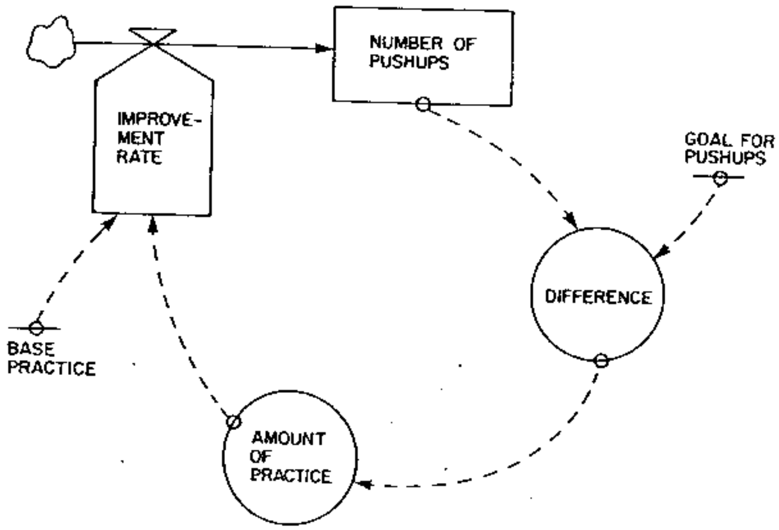

Figure 15.13 Flow diagram of pushups and practice, with goal

time to find out the exact relationship between the size of the "gap" and the amount he practices. But a plausible hypothesized relationship might be that Jim will practice one-half minute per day for each pushup he desires but cannot do:

Difference(pushups) = Goal for Pushups − Number of Pushups = 50 − Number of Pushups

Amount of Practice(minutes) = 0.5(minutes/pushup)*Difference (pushups)

Thus, for example, if Jim can do 20 pushups, while wanting to do 50, he practices (0.5)(50 − 20) = 15 minutes a day.

Now, all that remains is to specify the relationship between the amount Jim practices and the rate at which he improves. For simplicity, we might as well retain the assumption used in Example IV.

Improvement Rate(pushups/month) = (Amount of Practice − 10 minutes)*0.2(pushups/month/minute)

Exercise 19: Pushups and Practice Model, with Goal

a. Write DYNAMO equations for the pushup and practice model. (Choose DT = 1 month.)

b. Run the model setting the initial number of pushups Jim can do to 30. Rerun the model, setting the initial number he can do to 20. Rerun the model again, setting the initial number of pushups to 10. How do the results differ? Why?

c. How does the model's equilibrium point depend on the amount of time Jim practices per desired pushup he cannot do? Run the model several times, choosing different values for this parameter, Practice Time per Desired Pushup.

d. How does the equilibrium point depend on the value 0.2 in the equation relating the amount Jim practices and his improvement rate? Try running the model with alternate values for this parameter.

e. How does the equilibrium point depend on the value 10 in the equation relating the amount Jim practices with his improvement rate? Run the model with alternate values of this Practice Effectiveness parameter.

f. What changes in equation structures or parameters would enable Jim to reach his goal of 50 pushups? Run the model to demonstrate this.

CHOOSING A VALUE FOR DT

In the exercises and examples so far, an important technical issue has been treated lightly. The simulation technique used proceeds iteratively, stepping through time in intervals of length DT. In the library example, for instance, DT = one year; in the pushups example, DT = one month; in the yeast case, DT = one hour; and in the coffee cooling example, DT = one minute. Why were these values chosen to use in simulating the models? What would have happened had different values for DT been used?

The yeast example will be used to examine some of these issues more carefully. Table 15.1 shows four simulations of the yeast budding positive feedback loop, each using a different value of DT. (The four values are DT = 10, DT = 1, DT = 0.5, and DT = 0.1.) Table 15.2 shows four simulations of the yeast deaths negative feedback look, using the same four values of DT.

Table 15.1 Yeast budding simulations

Time (hours)	*DT = 10*	*DT = 1*	*DT = .5*	*DT = .1*
0	10.000	10.000	10.000	10.000
5		16.11	16.29	16.45
10	20.000	25.94	26.53	27.05
15		41.77	43.22	44.48
20	40.000	67.28	70.40	73.16
25		108.35	114.67	120.32
30	80.000	174.49	186.79	197.88

Table 15.2 Yeast deaths simulations

Time (hours)	DT = 10	DT = 1	DT = .5	DT = .1
0	10.000	10.000	10.000	10.000
5		7.738	7.763	7.783
10	5.000	5.987	6.027	6.058
15		4.633	4.679	4.715
20	2.500	3.585	3.632	3.670
25		2.774	2.820	2.856
30	1.250	2.146	2.189	2.223

As can be seen, in both the yeast budding and yeast deaths examples, the precise numerical results produced by the simulations differ depending on the values chosen for DT. For example, in the yeast budding case, the number of yeast cells grows most rapidly when DT = 0.1, and least rapidly when DT = 10. At first glance, this seems puzzling, since the budding fraction in all four cases is identical: BUDFR = 0.1. But a moment's reflection suggests an explanation: The situation is exactly analogous to compound interest. When DT = 1, for example, new cells are added to the yeast population exactly once per hour. When DT = .5 new cells are added every half hour, and thus new cells produced in one half-hour interval can themselves produce new yeast buds in the next half hour. (Hand-simulating the results for two or three hours, using DT = 1 and DT = 0.5, should illustrate this adequately.)

Which (if any) of these values for DT is correct? Certainly, yeast cells do not wait until the exact stroke of each hour to bud. For that matter, they surely do not bud exactly on the half hour, or at quarter-hour intervals. Presumably, individual yeast cells bud at various times throughout the hour.

For all practical purposes, it seems appropriate to assume that yeast cells bud more or less continuously. That is, at any moment, some yeast cells are in the process of budding. Similarly, at any moment, some are in the process of dying. But if this is true, how small a value of DT is needed to represent appropriately the yeast cell behavior?

It would be possible, of course, to simulate yeast budding using a DT of one minute, or even one second, if necessary. The main restriction is a practical one. The smaller the value selected for DT, the more computer time required to run the model (since more iterations are required). Thus choosing small values of DT increases the cost of running a model and lengthens the time spent sitting at the computer waiting for a model run to be completed.

In general, the proper approach in choosing DT is to select a value small enough to provide a reasonable approximation of the process being modeled, but not a value so small that it requires unnecessary computation time.

How small a value of DT is small enough? Consider the yeast deaths case first. Notice that when DT = 10, the number of yeast cells drops most quickly, falling to 1.25 cells in 30 hours. When DT = 1, the number of cells falls somewhat less rapidly, reaching about 2.1 in 30 hours. The results for DT = 0.5 and DT = 0.1 are very similar to the results for DT = 1. Thus in simulating the process of yeast deaths, it seems reasonable to choose DT = 1. Choosing a value of DT less than one would not produce a noticeably better approximation of the process of yeast deaths, and it would use up unnecessary computer time.

Now turn to the yeast budding loop. When DT = 10, the number of cells rises most slowly, reaching 80 cells in 30 hours. When DT = 1, the number of cells grows to about 174; when DT = 0.5, the number of cells reaches about 187; and when DT = 0.1, the number of cells rises to 197. The results for DT = 0.5 and DT = 0.1 are fairly similar to one another, although they still differ a bit. In simulating the process of yeast budding, a value of DT = 0.5 or possibly DT = 0.1 might be best. (The value used in Chapter 14, DT = 1, is perhaps a bit large.)

When formulating a model, it is important to give some attention to the choice of DT *before* using the model to draw any final conclusions about system behavior. The easiest way to choose an appropriate value of DT is to try various values, until one small enough is found, such that still smaller values do not produce noticeable changes in the simulated results. Once an appropriate choice of DT has been chosen, that value can, in general, continue to be used when employing the model to analyze system behavior and to test proposed policies. (Of course, experiments with DT should *not* be conducted when the model is in equilibrium, or nothing will happen!)

A few general rules of thumb can often provide a good starting point in selecting DT. Models that generate exponential growth require a DT that is much smaller than the doubling time involved. It is often a good first step to choose a DT from 1/5 to 1/10 the doubling time. (For example, the doubling time in the yeast budding case is roughly 7 hours. Thus a value of DT around one-half hour is a reasonably conservative choice.) Models that generate exponential growth over an extended period of time are especially sensitive to the choice of DT. (For example, the yeast population doubles roughly 4 times in a 30-hour simulation run, and thus, small errors mount up fairly quickly.)

Models involving negative loops require a DT smaller than the halving-times associated with the negative loops. A DT of roughly 1/3 or 1/4 the halving time is often a reasonable choice. (For example, the halving time in the yeast deaths case is roughly 15 hours. Thus any value of DT shorter than 4 or 5 hours is appropriate.)

One difficulty in trying to apply these rules of thumb is that, for complex models, it is often hard to determine, in advance of simulating the model, what the exact doubling times or halving times involved in various loops might be. In many cases, the rules of thumb provide only a rough general idea, and model experiments must be used to select an appropriate DT.[3]

Exercise 20: Experiments with DT

a. Run the coffee cooling model, trying various values of DT.

b. How does the choice of DT affect model behavior?

c. Which value of DT do you think is best?

Exercise 21: Jobs and Migration—Part II

Review Exercise 16 earlier in the chapter, and then do the following:

a. Use your judgment to select some plausible parameters for the model of jobs and migration you developed there. (You might want to choose parameters that reflect a hypothetical city.)

b. Write DYNAMO equations for your model.

c. Simulate the model, experimenting with various values of DT.

d. Determine the equilibrium level of population in the city.

e. Suppose the population in your city is in equilibrium, but in two years a major industrial company in the town closes, causing the number of jobs to fall by 10 percent. Use a STEP function to analyze the response of the system to the sudden decline in jobs.

ENDNOTES

1. Students with some background in calculus will recognize that the yeast budding model can be written in differential equation form as

$$\frac{dy}{dt} = by(t),$$

where $y(t)$ = the number of yeast cells at time t, and b = the budding fraction. The solution to the equation is $y(t) = y(0)e^{bt}$, where $y(0)$ = the initial number of yeast cells.

This solution to the differential equation can be used to derive the doubling time for the number of yeast cells. The number of cells will double when $y(t) = 2y(0)$, and this takes place when $e^{bt} = 2$. Taking logarithms, the number of cells will double when $bt = ln2$. Since $ln2$ is roughly 0.7, and the budding fraction $b = 0.1$, the doubling time $t = (0.7/0.1) = 7$ hours.

2. The differential equation for the yeast deaths loop can be written $dy/dt = -y(t)/a$, where $y(t)$ = the number of yeast cells at time t, and a = the average lifetime of yeast. The solution to this equation is $y(t) = y(0)e^{-t/a}$, where $y(0)$ = the initial number of yeast cells.

This solution can be used to derive the halving time for the number of yeast cells. The number of cells will fall to half its initial value when $y(t) = ½y(0)$, and this takes place when $e^{-t/a} = ½$. Thus taking logarithms, the number of cells will reach one-half its initial value when $-t/a = ln(½)$. Since $ln(½)$ is roughly equal to -0.7, and the average lifetime of yeast $a = 20$, the halving time $t = (20)(0.7) = 14$ hours.

3. In discussing the choice of DT, it has been assumed that the process being modeled takes place continuously in time. In the yeast case, for example, at any moment, some yeast cells are in the process of budding and others are in the process of dying. The coffee cooling case is a particularly good example of a continuous process. For an actual cup of coffee, a certain amount of cooling occurs every second, or every micro-second for that matter.

   In some cases, the assumption that the process being modeled is continuous may seem less valid. For instance, in the pushup case, Jim might practice exactly once each day, at 2:00 in the afternoon. Thus at first glance, it may seem that any improvement that takes place in the number of pushups Jim can do is not continuous, but occurs in once-a-day jumps. It is nevertheless possible that some of the improvement in Jim's ability to do pushups occurs following his practice session each day, while he is eating, sleeping, and so forth. Although the daily practice sessions are not continuous, the overall process of improvement can be viewed as at least approximately continuous.

   By and large, for most social and economic systems, it is reasonable to assume that the process being modeled is roughly continuous, at least to a first approximation. In cases where the processes being modeled are quite clearly not continuous, the methods of analysis described in this chapter are not strictly appropriate.

Flow diagram symbols


Level

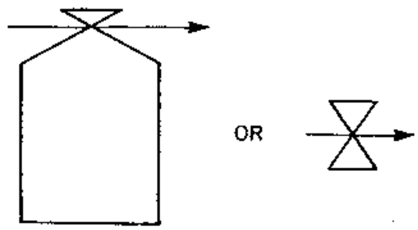

Rate

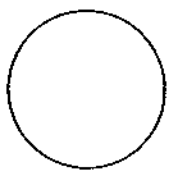
Auxiliary

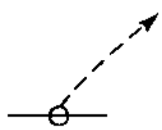
Constant

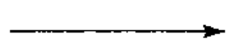
Flow

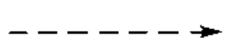
Cause-and-Effect Link


Source or Sink