Chapter 4

Water Flows in the Mono Basin

Mono Lake is an ancient inland sea in the eastern Sierra. It is one of the oldest continuously existing lakes in the world, and the land forms reflect the lake's ancient history. Volcanic islands rise in the middle of the lake, and tufa towers rise around the edges. The basin is a land of stark contrasts and spectacular vistas. It's also an oasis for wildlife in the high desert country. Migratory birds use the lake as a stopover location; nesting birds raise their chicks on the islands. The birds are drawn by a simple but extraordinary ecosystem. Microscopic algae thrive in the lake, providing the food supply for brine shrimp and brine flies. These are astoundingly prolific organisms that can provide a virtually limitless food supply for birds under proper conditions.

Mono Lake was selected for your first experience with a "real-world" model because it provides important lessons for policy making as well as simulation modeling. From a policy point of view, Mono Lake is a story of how

> a handful of people began a campaign to save a dying lake, taking on not only the City of Los Angeles, but the entire state government by challenging the way we think about water. Their fight seemed doomed in the beginning, but long years of grassroots education and effort finally paid off in 1994, when the California Water Resources Control Board ruled that Los Angeles's use of Mono Basin waters be restricted. Over time, the lake will return to a healthy condition. . . . The battle over Mono Lake is one of the longest and most fiercely contested conservation battles in US history, and that rare one with a happy ending. (Hart 1996)

Mono Lake is well suited to demonstrate the power of stock-and-flow modeling. By the end of this chapter, you will see a model that may be used to simulate changes in the lake level with different policies controlling water export to Los Angeles. You may then expand the model and test your own policies for controlling the size of the lake.

Background

The Mono Lake story began early in the century when the City of Los Angeles looked to the Owens Valley for new sources of water. Under the direction of William Mulholland, the city completed the Owens Valley aqueduct in 1913. Los Angeles was a city of 100,000 in the year 1900. By the year 1930, the population had reached 1 million, and the city was looking

Figure 4.1. Location of Mono Lake and the aqueducts serving Los Angeles.

beyond Owens Valley. By 1941, the Colorado River Aqueduct was completed. And in the same year, the Los Angeles Aqueduct was extended to reach the Mono Basin. Figure 4.1 shows Mono Lake and the aqueducts that serve Los Angeles.

The city began diverting water from the Mono Basin in 1941. Stream flows toward the lake were diverted into a tunnel running beneath the Mono Craters to reach the northern Owens River. The journey to Los Angeles is nearly four hundred miles, and the water flows by gravity and siphons the entire way, producing hydroelectric energy en route. By the 1970s, diversions averaged around 100 KAF/yr (thousand acre-feet per year).

The impact of the diversions is evident from the chart in figure 4.2. The lake's surface was measured at 6,417 feet above sea level in 1941. The lake held around 4.3 million acre-feet of water, and its surface area spread across 55 thousand acres. Salinity, a crucial factor for the brine flies and brine shrimp, was around 55 g/L (grams per liter). Then the diversions began. Figure 4.2 shows the steady decline in the lake's elevation during the next four decades. By

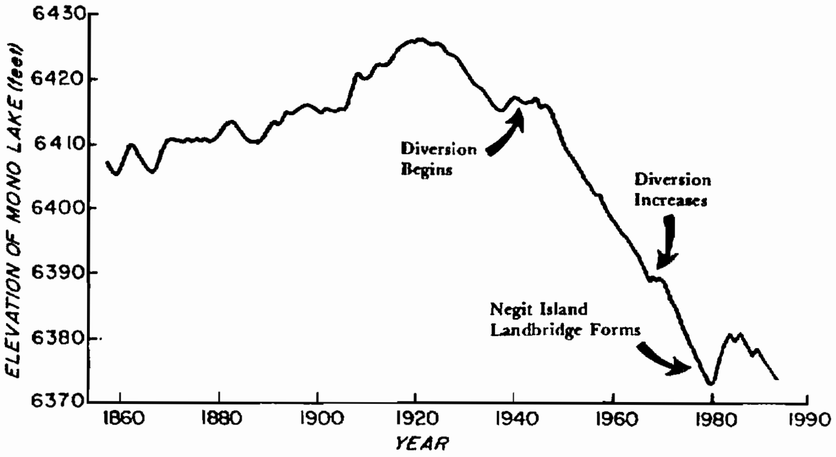

Figure 4.2. Historical elevations of Mono Lake.

Photo 4.1. The Northwest corner of Mono Lake with the land bridge connected to Negit Island. Source: Mono Lake Committee Collection.

the year 1981, the lake's volume had been cut approximately in half, and its salinity had climbed to around 100 g/L. The lake stood at 6,372 feet, 45 feet below its position when diversions began. As the lake shrinks, salinity climbs, and higher salinity can reduce algae production and lower the survivability of brine flies and brine shrimp. When these herbivores decline in number, the nesting birds may not find adequate food to raise their chicks. The migrating birds may not be able to add sufficient weight for the next leg of their migration. A declining lake level poses other dangers as well. When the lake receded to 6,375 feet, for example, a land bridge was formed to Negit Island (see photo 4.1), and a once secure nesting habitat became vulnerable to predators.

Reversing the Course

Environmental science students studied the lake during the 1970s. They were alarmed at what they found and fearful for the future of the lake. For example, they feared that higher salinity could lead to serious declines in brine shrimp population and a subsequent loss of suitable habitat for the bird populations. In 1978, one group of students formed the Mono Lake Committee, a grassroots education and advocacy group. Mono Lake also drew the attention of the National Audubon Society, which filed suit against the City of Los Angeles in 1979. The California Supreme Court responded in 1983. It held that the public trust mandated reconsideration of the city's water rights in the Mono Basin. The court noted that Mono Lake is a scenic and ecological treasure of national significance and that the lake's value was diminished by a receding water level. The court issued an injunction later in the 1980s limiting the city's diversions while the State Water Resources Control Board reviewed the city's water rights.

The Control Board considered a variety of alternatives for the future. One extreme was the "no restriction" alternative, in which the city would be free to divert water as in the past. With no restrictions, the Control Board expected the lake to decline for another fifty to one hundred years and reach a dynamic equilibrium at around 6,355 feet (Jones and Stokes Associates 1993, 2–17). The opposite extreme was the "no diversion" alternative. If all of Mono Basin's streams were allowed to flow uninterrupted to the lake, the Control Board expected the lake to climb over a period of one hundred years, eventually reaching dynamic equilibrium at around 6,425 feet.

The Control Board issued its decision in 1994. It concluded that the appropriate balance between the city's water rights and the public trust would be served by allowing the lake to rebuild to a higher elevation. The target elevation is marked with a small plaque next to the boardwalk leading to the shore of the lake. The plaque explains that the lake stood at 6,392 feet back in the year 1963. The lake will now be allowed to rebuild toward that level.

Water Flows in the Basin

The preceding background provides the context for the modeling exercises in this chapter. If you wish to learn more, turn to the book's home page. It includes information on the Mono Basin and how the diversions fit within the water supply system for the City of Los Angeles. It also includes additional background, photographs, and links to related home pages. Figure 4.3 shows one of the sketches to be found on the home page. It depicts water flows in the basin in an average year. These flows will help us appreciate why the lake has declined historically.

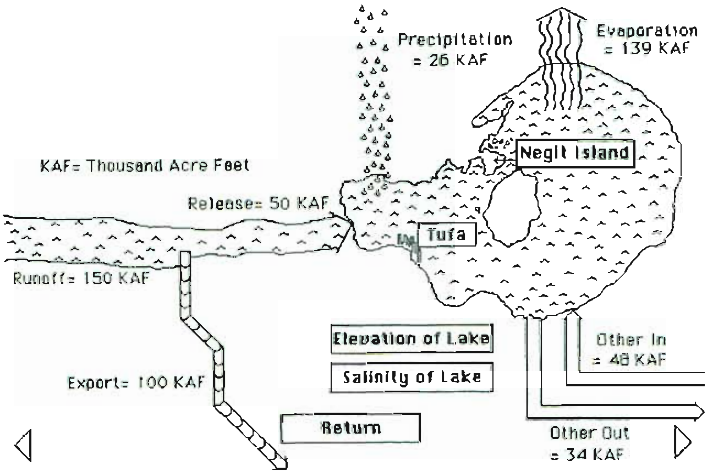

Figure 4.3. Illustrative flows in an average year.

The principal flow into the basin is the runoff from five streams that drain the Sierra Nevada. (Photo 4.2 shows Rush Creek, the largest of the five streams, during high runoff.) Vorster (1985) reviewed measurements of the five streams over the years 1937–83 and concluded that they amount to 150 KAF/yr with average weather. The City of Los Angeles operates diversion points on the gauged streams, and figure 4.3 depicts diversions of 100 KAF/yr. This leaves 50 KAF/yr to reach the lake. Precipitation and evaporation can vary from year to year depending on weather and the size of the lake. Evaporation is, by far, the largest flow out of the basin. The "Other Out" is a combination of several smaller flows such as evapotranspiration and exposed lake bottom evaporation, The "Other In" is a combination of several flows such as ungauged sierra runoff and municipal diversions. The total of all inflows in figure 4.3 is 224 KAF/yr. The outflows total 273 KAF/yr. The net result is a loss of around 50 KAF/yr.

Now imagine what would happen if this annual loss were sustained year after year. There

Photo 4.2. This photo shows flow in Rush Creek, one of the five gauged streams that drain the Sierra Nevada and flow toward Mono Lake. Source: Mono Lake Committee Collection.

would be less and less water stored in the lake. After twenty years, the lake would have lost 1 million acre-feet. If the loss continued for forty years, the lake would lose 2 million acre-feet. This simple arithmetic is sufficient to illustrate how a lake with 4.3 million acre-feet could be cut approximately in half with forty years of diversions. Let's turn now to computer simulation modeling to look at this problem in more detail.

Purpose of the Model

It is crucial at the outset of any project to specify the purpose of the model as clearly as possible. Our purpose is to project the future size of the lake given different assumptions about the amount of water exported to Los Angeles. We expect to use the model to make projections that resemble the graph shown in figure 4.4. This initial graph is called the *reference mode.* It serves as a target pattern of behavior that you expect the model to generate. The reference mode is normally drawn based on our understanding of the fundamental patterns of growth, decay, overshoot, or oscillation explained in chapter 1. In each new problem area, we should ask ourselves if our dynamic problem corresponds to one of the standard shapes. When classifying the Mono Lake problem, the best shape is decay. If there were no influx of water each year, Mono Lake would decay to zero. But nature provides precipitation and runoff each year, so our intuition tells us that the lake will probably decay to a smaller size, but not all the way to zero.

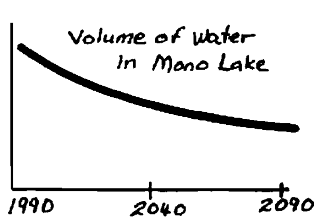

Figure 4.4. The reference mode for the Mono Lake model.

The reference mode begins in the year 1990. (This is the year when one of my students did a class project on the lake, but you may equally well imagine starting the analysis in the current year.) Figure 4.4 shows the reference mode extending from the year 1990 to the year 2090. This long time period seems reasonable given the Control Board's assessment of the time needed for the lake to reach dynamic equilibrium. The length of time appearing on the horizontal axis of the reference mode is often called the *time horizon* of the model.

With such a long time horizon, we can now begin to make useful decisions on what to include and what to exclude from the model. When looking forward one hundred years, let's ignore daily or monthly variations in the lake volume. Notice that figure 4.4 is a smooth line that shows a gradual decline in the volume of water in the lake. It makes it clear that we are ignoring seasonal variations. The shape of the reference mode also reminds us that we are ignoring year-to-year variations in the weather. One year may be particularly wet; the next, particularly dry. These variations could be important in a short-term model, but they can be ignored in a long-term model. For now, let's assume that every year is an average year as far as the weather is concerned.

One final feature of figure 4.4 should catch your attention—there are no numbers on the vertical axis. You might have wondered why we don't show the lake volume at around 2 million acre-feet at the start and show estimates of the final volume in the year 2090. You should feel free to include or exclude numerical detail when sketching the reference mode. If the numbers are available and they improve communication among members of the team, include them in the sketch. If not, feel free to leave them out. It's the general shape of the diagram that counts at this stage of the process.

Drawing a reference mode before building the model may seem like "rigging" the analysis. Some would say that you should wait until the model is done and then see how it behaves, Don't let this thinking divert you from this first step. Think about the process that you went through the first time you built a model airplane. You didn't build the model blind, thinking that you would learn its likely pattern of behavior once you threw it into the air. You built it with a predetermined image of the behavior. You probably envisioned that the model could glide through the air for five to ten feet, maybe longer if you were clever with the design. Building a system dynamics model is similar; you start with a predetermined image of the model's likely behavior over time. Drawing a reference mode is key to the dozens of decisions you must make later in the model building process. Your choice of a reference mode will guide which factors are included (and which factors are excluded) from the model.

List the Policies and Start Simple

Specifying policies at the outset is also important because it guides our decisions about what to exclude from the model. For our purposes, there is one key policy variable: the amount of water exported from the Mono Basin. Our purpose is to learn how the size of the lake varies with variations in water exports. In this example, we must be sure that *exports* are included explicitly in the model. Other policy variables (like the demand for water or the price of water in Los Angeles) are not on our list, so they may be excluded from the model.

Building a model is an iterative, trial-and-error process. The best approach is to begin with the simplest possible model that could explain the reference mode. You should build, test, and reflect on the first model before moving to a more complicated model. The worst thing you can do is try to build the "perfect model" right from the start. You will make more progress in the long run if you start simple and learn as much as you can from each new model. Now, before reading ahead, ask yourself what you would do next. What are the stock variables for a model of Mono Lake, and what are the flows that directly influence those stocks? What units would you use to measure the stocks and flows? Then think of the converters you would use to fill out the picture.

An Initial Model of Mono Lake

Figure 4.5 shows a first cut at a model. A single stock is used to keep track of the volume of *water in lake*. The volume of water has the feel of a stock variable because it changes more slowly than other variables. If there was zero runoff this year, for example, there would still be some water in the lake next year. Also, the water in the lake seems to make sense as the main place where accumulation takes place in the system. Let's measure the stock in KAF (thousand acre-feet). We may set the initial value of the stock at 2,228 KAF to represent the water stored in Mono Lake in 1990. The model is specified to run in years, so the units for each of the flow variables must be KAF/year. To keep the units consistent, the area of the lake will be measured in thousands of acres, and the annual rates of precipitation and evaporation will be measured in feet per year.

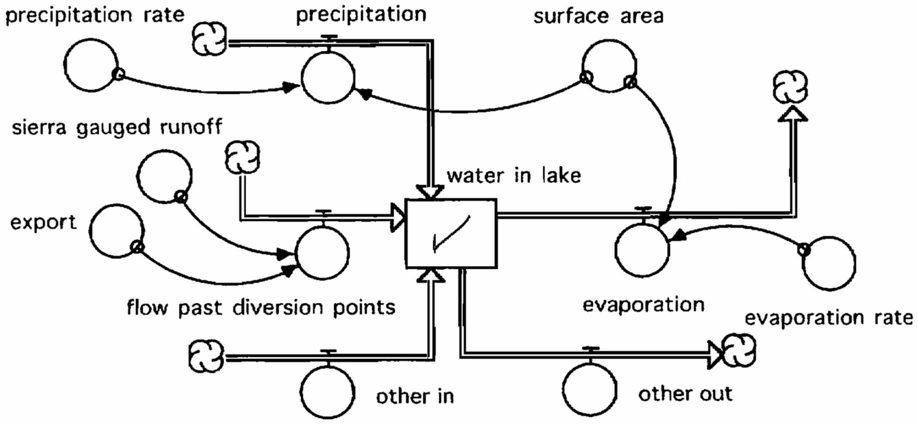

Figure 4.5. The first model of Mono Lake.

You've seen in figure 4.3 that the main flow into the basin is the runoff from the sierra streams. To match the names used by Vorster (1985), let's name this converter *sierra gauged runoff.* The *export* stands for the amount of water diverted into the tunnel to leave the basin. We would subtract export from the gauged runoff to get the *flow past diversion points,* the flow that actually reaches Mono Lake. The model includes *evaporation* and *precipitation* flows as well. The *evaporation* is the lake's *surface area* multiplied by the *evaporation rate,* while the *precipitation* is the surface *area* multiplied by the *precipitation rate.* Let's set the area at 39 thousand acres based on the area in 1990, and let's assume that the evaporation rate is 3.75 feet/yr and the precipitation rate is 0.67 feet/yr in an average year. Other flows are the collection of *other in* flows and a collection of *other out* flows. The numerical values will be set at the average year values shown previously. The equations for this first model are listed in figure 4.6.

Now, what do you expect to see from this initial model? Will it generate the reference mode? Figure 4.7 shows the simulation result. Flow past the diversion points is constant at 50 KAF/yr. The evaporation is at around 146 KAF/yr. The water in the lake declines in a linear fashion throughout the simulation. Indeed, the volume reaches zero by the year 2030, and it

```
water_in_lake(t) = water_in_lake(t - dt) + (flow_past_diversion_points +
other_in + precipitation - other_out - evaporation) * dt
INIT water_in_lake = 2228
INFLOWS:
flow_past_diversion_points = sierra_gauged_runoff-export
other_in = 47.6
precipitation = surface_area*precipitation_rate
OUTFLOWS:
other_out = 33.6
evaporation = surface_area*evaporation_rate

evaporation_rate = 3.75
export = 100
precipitation_rate = 0.667
sierra_gauged_runoff = 150
surface_area = 39
```

Figure 4.6. Equations for the first model.

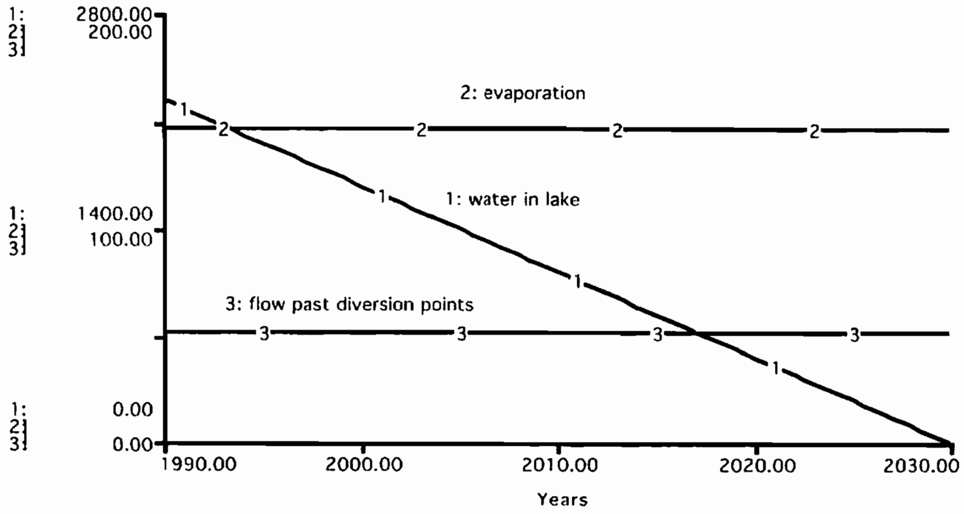

Figure 4.7. Simulation results from the first model of Mono Lake.

would continue to decline if we were to extend the simulation further. These results are clearly spurious. Not only do they not match the reference mode, but it makes no sense to see the *water in lake* become negative.

At this point, you might be tempted to invoke the "non-negative" option on the stock variable. But this would simply be covering up a fundamental problem with the model. Also, you might wonder if the spurious behavior is the result of the particular parameters. Perhaps we should try a simulation with lower *export* or a simulation with a different estimate of the *surface area* of the lake. Do you think these experiments would change the linear pattern? Would they ever allow the model to generate the reference mode?

You can give these experiments a try. You'll soon discover that this model will never generate the target pattern. No matter what set of input parameters you try, the pattern will turn out to be either linear decline or linear increase. The problem with this first model is not the input parameters; it's the structure. It's time to expand the stock-and-flow diagram to improve the structure.

Second Model of Mono Lake

If Mono Lake were shaped as a cylinder, its surface area would remain constant as the volume falls. But the lake looks much more like a shallow cup than a cylinder. Its surface area tends to shrink as the volume falls. For the second model, let's change the *surface area* to an internal variable that will decline with a decline in the volume of water in the lake. To introduce some new terminology, one might say that we have changed the area from an ***exogenous*** variable to an ***endogenous*** variable. The challenge is to represent the *surface area* as a function of the volume of the lake. Perhaps we could write the equation for the surface area of a cone as a function of the volume of water stored in the cone. This approach is not likely to be realistic, however, because the shape of the lake bottom is complicated due to islands and volcanic structures. Fortunately, geologic and bathymetric surveys have been completed to provide all the detail we need. The survey results are reported by Vorster (1985, 261) and are summarized in table 4.1 in terms of volume, area, and elevation of the lake.

Table 4.1. Survey results.

Volume of water in lake (KAF)	*Surface area (Kacres)*	*Elevation (feet above sea level)*
0	0	6,224
1,000	24.7	6,335
2,000	35.3	6,369
3,000	48.6	6,392
4,000	54.3	6,412
5,000	57.2	6,430
6,000	61.6	6,447
7,000	66.0	6,463
8,000	69.8	6,477

To incorporate the survey information in the model, we would use an information connector from the *water in lake* to the *surface area* as shown in figure 4.8. Then invoke "become a graph"; select *water in lake* to appear on the horizontal axis and ask for 9 points on the horizontal axis. Set the lower bound of the horizontal axis at 0 and the upper bound at 8,000; then enter the nine values of surface area from the survey table.

The results from the second model are shown in figure 4.9. They are certainly different from those of the previous model. The *surface area* begins the simulation at 39 thousand acres

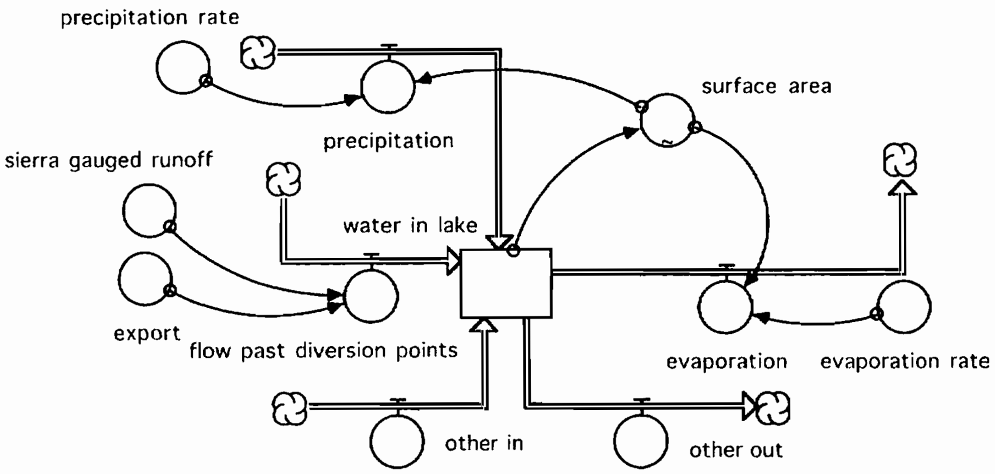

Figure 4.8. The second model of Mono Lake.

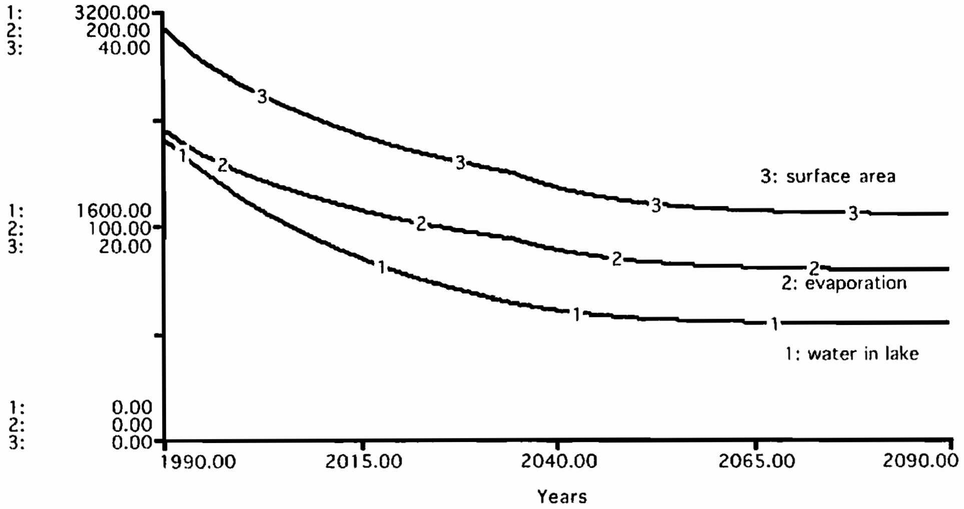

Figure 4.9. Simulation results from the second model of Mono Lake with export held constant at 100 KAF/year.

and declines as the volume of water declines. *Evaporation* is around 146 KAF/yr at the start of the simulation, but it declines over time as the lake shrinks in size. The pattern in figure 4.9 is just what we are looking for. We see a gradual decline to dynamic equilibrium.

This is good progress; we have a relatively simple model that generates the reference mode. But a serious limitation of this model is that it provides no indication of whether the lake's ecosystem is in danger. What can be done to represent the state of the ecosystem without having to simulate the complexities of the food web? One approach is the "proxy approach." We search for a variable that would be easy to include and allow the new variable to serve as a proxy for the state of the ecosystem. Most discussions of the threats to Mono Lake's ecosystem rely on elevation as a proxy. For example, several investigators have cautioned that we should be watchful for different categories of problems if the lake falls to the following elevations:

6,380 feet—severe dust storms

6,375 feet—general ecosystem decline

6,372 feet—major loss of gull nesting habitat

6,363 feet—critical salinity levels

6,352 feet—general collapse of the ecosystem

For the next iteration, let's add elevation to the model. And while we are making changes, let's elaborate on the description of some of the flows.

A Third Model: Elevation and Flow Details

Since the elevation is reported in the survey data, we may add elevation as a converter that is linked to the volume with a connector. Then invoke the "become a graph" feature with the volume of water on the horizontal axis. Design the axis for nine entries and enter the nine values of elevation from the survey table. Figure 4.10 shows the new diagram. The other changes appear at the bottom of the diagram where we have more detail on the other flows in and out of the basin. *Other out* is now the sum of four separate flows, each of which has been estimated by Vorster (1985) as follows:

- Net evaporation from Grant Lake, a small lake upstream from Mono Lake, is 1.3 KAF/yr.
- Evapotranspiration from irrigated land, riparian vegetation, and vegetation such as salt grass, greasewood, and willows is 13 KAF/yr.
- Exposed lake bottom evaporation is estimated at 12 KAF/yr. (This refers to the evaporation from residual pools of water stranded by the receding lake as well as evaporation from groundwater brought to the surface by capillary action.)
- Some groundwater is intercepted by the underground conduit and exported south to Los Angeles. Vorster estimates "groundwater export" at 7.3 KAF/yr.

The *other in* flow is now calculated as the sum of four separate flows, each of which has been estimated by Vorster (1985) as follows:

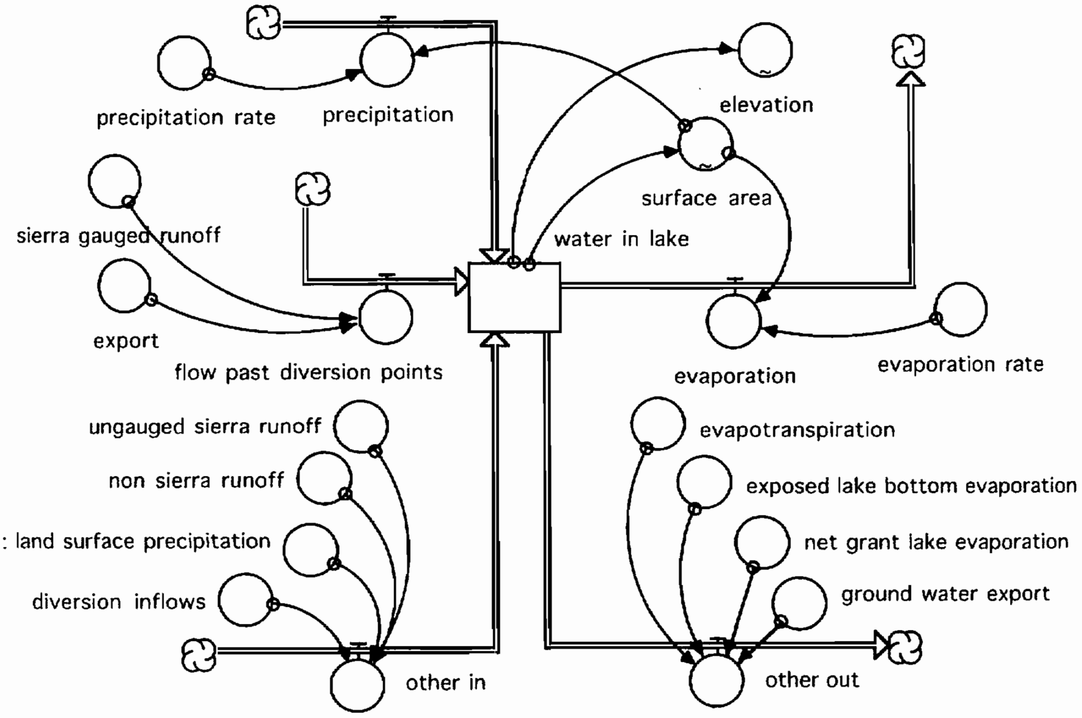

Figure 4.10. The third model of Mono Lake.

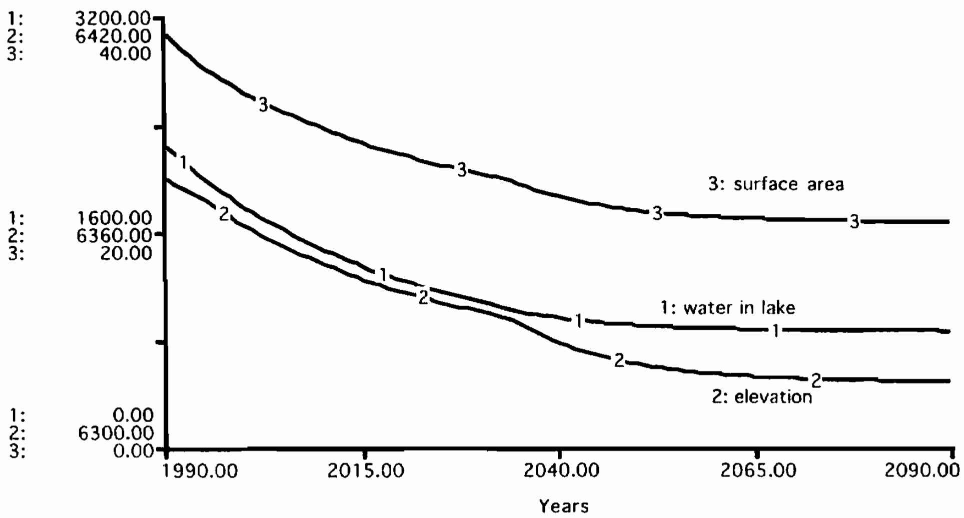

Figure 4.11. Results from the third model.

- Ungauged sierra runoff from small and intermittent watersheds is expected to be 17 KAF/yr.
- Nonsierra runoff from watersheds to the north, east, and south are estimated to yield 20 KAF/yr based on a soil moisture surplus method.
- Vorster refers to precipitation falling on the basin (but not on the lake) as "net land surface precipitation," and he estimates it at 9 KAF/yr.
- The small municipalities divert water from Virginia Creek (outside the basin) and allow the used water to flow toward Mono Lake, creating what Vorster calls "diversion inflows" of 1.6 KAF/yr.

You can double-check the sums to ensure that the new model has the same flows as before. We are not expecting the new model to show different behavior. Rather, we are looking for more clarity in communicating the nature of the flows.

The results of the third model are shown in figure 4.11. The water in the lake follows the same downward path as before. *Elevation* is the new variable, and we see that it follows an irregular downward path, especially around the year 2040. The lake has a complex shape, and there is a so-called nick point at around 6,368 feet where the topographic relief changes markedly. But the irregular pattern of *elevation* in figure 4.11 is probably not showing us the "nick point." The irregularities are probably due to the crude representation of the elevation with only nine survey points. By the end of the simulation, Mono Lake has declined to under a million acre-feet of volume, and its elevation is below 6,320 feet.

Checking the Model

One way to check the model is to set the initial value of the stock of water in the lake to match the volume back in 1941. We could then set the model inputs to match time series information during the four decades of diversions. Our confidence in the model would be bolstered if it could simulate the 45-foot decline in lake elevation that occurred during the interval from 1941 to 1981. Another test would be to check year-by-year variations in the simulated size of the lake with recorded variations. For example, if we simulated the unusually wet conditions in 1983, we would look to see if the simulated lake rose by 6 feet. We would also check the model by comparing it with more detailed models of the lake. One of the best is Vorster's (1985) water balance model. Vorster's thesis is exemplary because of the careful documentation, the independent estimates of each flow, and the historical comparisons to check accuracy. Since Vorster's model has been carefully checked against historical records, we might proceed on the assumption that it provides a good benchmark for the future. Furthermore, many of the parameter values in the third model are taken from Vorster's documentation, so we should expect to see similar results when the models are simulated under similar hydrologic scenarios. Vorster calculates the dynamic equilibrium of the lake under average climatic conditions with differing values of export. For a constant export of 100 KAF/yr, he expects the lake to equilibrate at 6,335 feet. This result is more than 15 feet higher than the elevation at the end of the simulation in figure 4.11. Something is missing from the model.

Fourth Model: Changing the Evaporation Rate

The missing factor is the change in the rate of evaporation as Mono Lake's water becomes more and more dense over time. The increased density arises from a fixed amount of dissolved solids held in solution in a shrinking volume of water. Highly saline waters tend to evaporate more slowly than fresh water due to a reduction in the vapor pressure difference between the surface of the water and the overlying air. But we don't need to add vapor pressures to the model to represent the change in evaporation. Rather, we can take advantage of evaporation studies documented by Vorster (1985, 90) and tabulated in table 4.2. This table uses specific gravity to measure the density of the water. (A value of 1.0 corresponds to fresh water; a value of 1.1 means the lake's water is 10 percent heavier than fresh water.) The second column reports a "multiplier" to summarize the impact of higher salinity. If Mono Lake's water is 10 percent heavier than fresh water, for example, its evaporation rate would be 92.6 percent of the evaporation rate for fresh water.

Figure 4.12 shows the flow diagram of a new model that includes the missing factor. The specific gravity depends on the *water in lake*, the *total dissolved solids*, and the *mass of fresh water*. The *mass of fresh water* is 1.359 million tons per KAF; the *total dissolved solids* is 230 million tons. The specific gravity is the mass of the actual water (with its dissolved solids) divided by the mass of the same volume of fresh water. For a lake with 2,228 KAF of *water in lake*, you should verify that the specific gravity would be 1.076. In other words, the water in the lake is 7.6 percent heavier than fresh water.

Table 4.2. Summary of evaporation studies.

Specific gravity of water in the lake	*Evaporation rate multiplier*
1.00	1.000
1.05	0.963
1.10	0.926
1.15	0.880
1.20	0.883
1.25	0.785
1.30	0.737
1.35	0.688
1.40	0.640

The *specific gravity* is then used to determine the *evaporation rate multiplier from specific gravity*, and the multiplier is used to find the actual evaporation rate. The multiplier takes on the value 1.0 when the specific gravity is at 1.0. Under this benchmark condition, the model would multiply the *fresh water evaporation rate* by 1.0 to get the actual *evaporation rate*. As the *specific gravity* increases above 1.0, however, the multiplier will decline below 1.0. If the *specific gravity* increases to 1.076, for example, the multiplier will turn out to be 0.94. In other words, the evaporation rate is 6 percent slower than the *fresh water evaporation rate*. The equation for the evaporation rate would multiply 0.94 times the *fresh water evaporation rate* of 3.75 feet/year to obtain the actual *evaporation rate* of 3.54 feet/year. Then, as in previous models, the *evaporation rate* is multiplied by the *surface area* of the lake to obtain the *evaporation*.

Figure 4.13 shows the simulation results of the new model with export held constant at 100 KAF/yr. The time graph shows volume, area, and elevation, the same measures shown before. It also shows the *specific gravity*. It begins the simulation at around 1.076, as you would

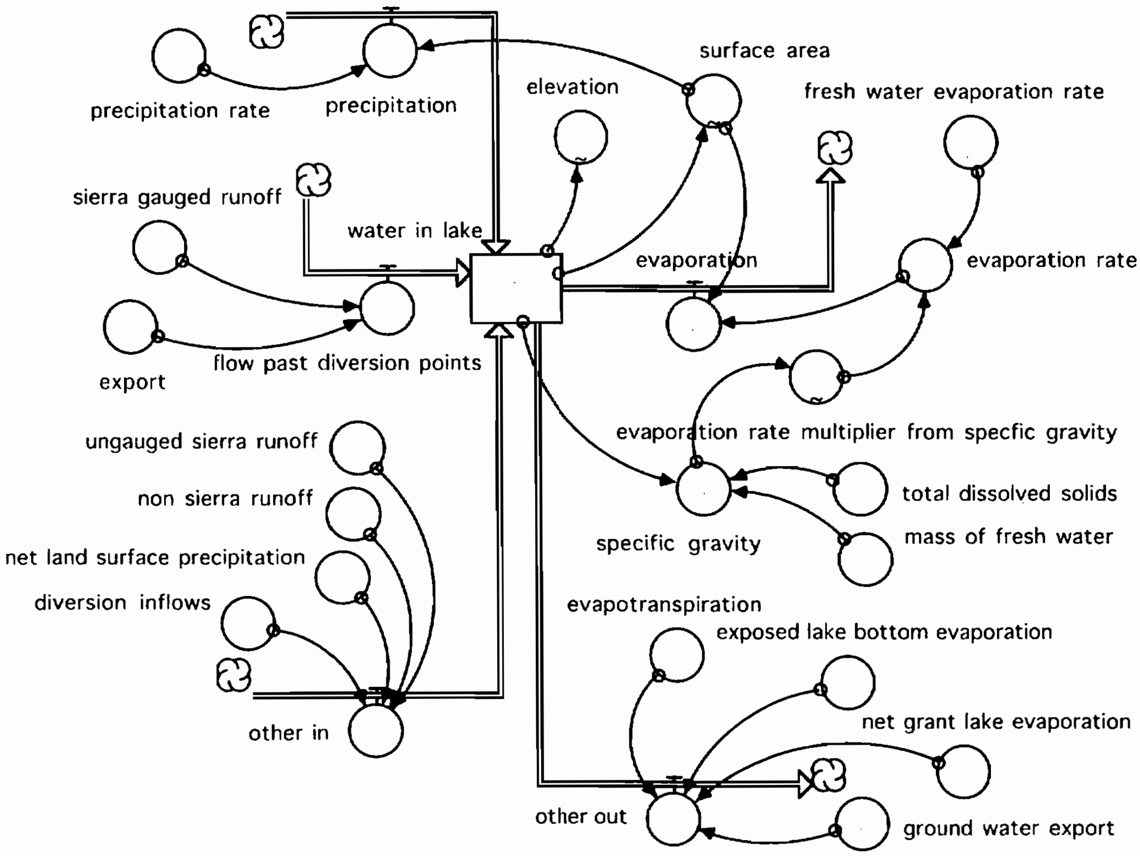

Figure 4.12. Stella flow diagram of the fourth model of Mono Lake.

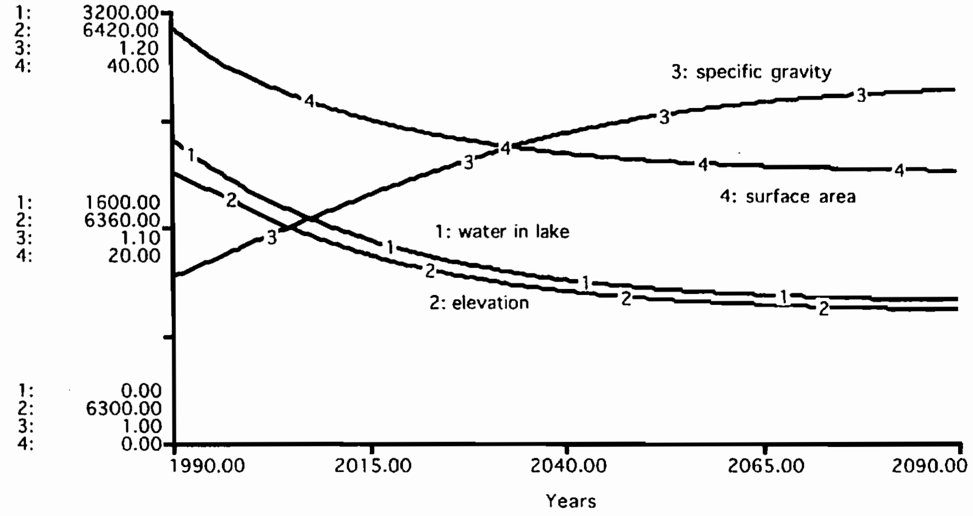

Figure 4.13. Simulation results from the fourth model of Mono Lake.

expect. It ends the simulation at 1.163. By the end of the simulation, the lake has declined to an elevation of 6,336 feet. The general pattern is the same as before—a gradual decline, eventually reaching dynamic equilibrium. And it's gratifying to see the equilibrium value within 1 foot of the corresponding result by Vorster (1985, 225).

Testing the Model

Now let's see what we might learn by testing the model. You've learned previously that system dynamics simulations should be used in pairs (like the blades of a pair of scissors). It takes at least two simulations to arrive at useful conclusions, so let's compare some simulations. For the first test, let's imagine that we allow the 100 KAF/yr of export to continue for the first half of the simulation period. Then we'll cut the export to zero for the second half. This is a crude but simple test to learn how responsive the lake is to a change in exports. Since water export was designated as a "policy variable" at the start of the chapter, this simple test might be called a "policy test" of the model.

Figure 4.14 shows the results of this first policy test. Elevation follows the same downward trajectory seen previously for the first half of the simulation. By the year 2040, the lake stands at 6,341 feet. After that year, the lake grows in size, eventually reaching an elevation of 6,398 feet. This test reveals that the slope of the recovery is somewhat steeper than the downward slope during the first half of the simulation. Because of the steeper recovery, the lake requires only around fifteen years to recover to the starting elevation. Then the recovery slows somewhat as the lake becomes larger and larger. This slowdown in the recovery is caused by the growing surface area of the lake. As the area increases, the lake's evaporation increases, and the lake's recovery is slowed. (And to a lesser extent, the slowdown in the recovery arises from the reduced specific gravity of the water in the lake.)

The policy test in figure 4.14 is meant to illustrate the type of tests you should conduct when first testing a model. It is simple and easy to interpret. But you should understand that this simple policy does not match the specific policies proposed or adopted in the Mono Basin. (These policy tests are left as exercises at the end of the chapter.)

Let's conclude the testing with an example of multiple simulations to portray the impact of a change in water export policy from the start of the simulation. Figure 4.15 shows the results of a sensitivity analysis of the importance of water export. This is a comparative graph showing the *elevation* from six different simulations of the model. Each simulation assumes a constant export from the beginning to the end of the simulation. The first run sets the export to zero, and the *elevation* climbs slowly over time. It reaches 6,410 feet by the year 2090 and is still growing at the end of the simulation. The sixth run sets export to 100 KAF/yr, and we see the same pattern as shown before. Two of the intermediate tests show the export that would leave the lake's elevation pretty much unchanged over the time interval. Exports are set to 40 KAF/yr in the third run and 60 KAF/yr in the fourth run, and these simulations suggest that the lake would remain relatively constant at the 1990 position. These simulations suggest that export must be cut approximately in half to preserve the lake at the 1990 level; they must be cut further to rebuild the lake to a higher elevation.

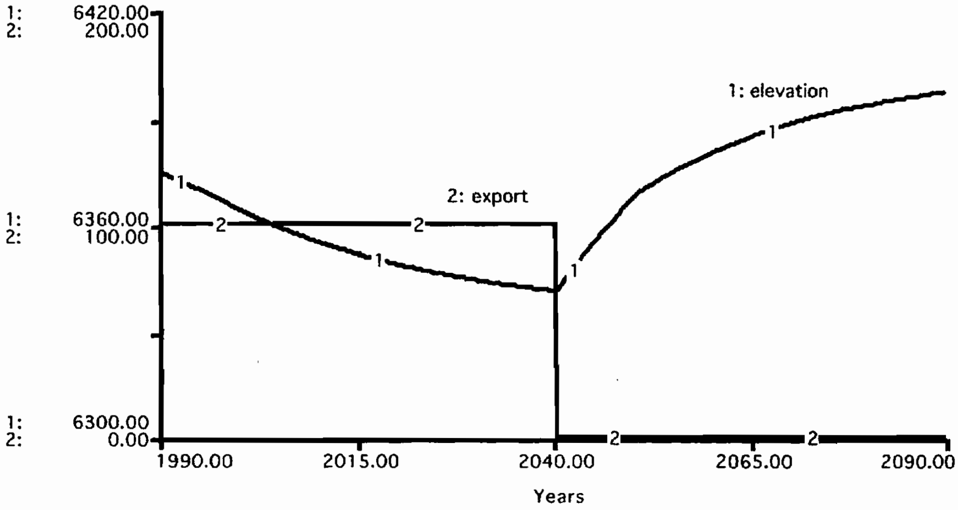

Figure 4.14. Simulated recovery after a change in export midway through the simulation.

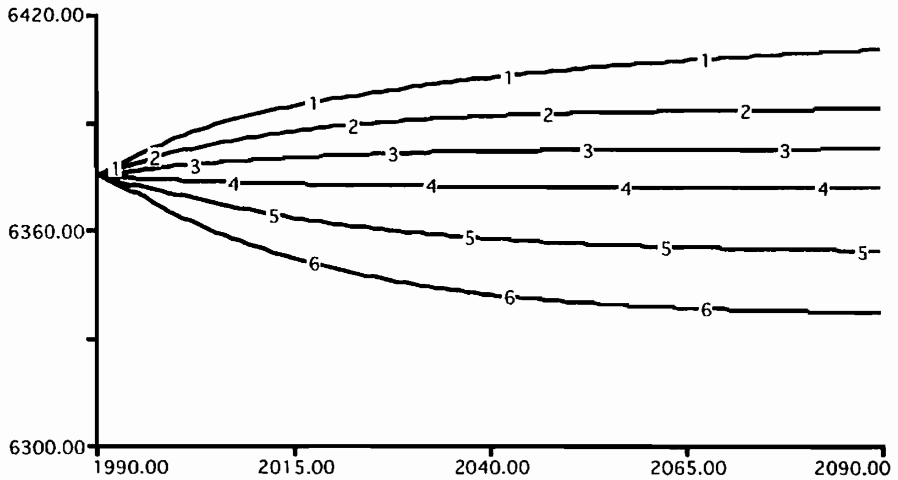

Figure 4.15. Simulated elevation of Mono Lake in a sensitivity analysis with export ranging from 0 (run 1) to 100 KAF/yr (run 6).

Conclusion

The fourth model is an excellent model of the water flows in the Mono Basin. Its main advantages are simplicity and clarity. Take a close look at the flow diagram in figure 4.12 and think of this way of describing a computer simulation model. This small diagram is sufficient to show every variable and every interconnection. The names are simple, and the units are clear. The simulations show the reference mode, and the results have been checked against a more detailed model. We now have a model that may be used to test the impact of different export policies.

Now, what do you think are the main problems with this model? Are you worried about the "average year" assumption? the lack of seasonal variation? the lack of precision in describing the geometry of the lake? Does it bother you that all five of the gauged sierra streams are combined into a single flow? These are some of the potential problems that may be addressed in student projects to improve the mode. As you add "improvements" to the model, use the policy results from the original model as a guide to whether your improvements are important. If your new model shows essentially the same policy results as the previous model, then you know your improvement is important only if it adds clarity and bolsters confidence in the model.

Perhaps the principal limitation of the Mono Lake model is that it is limited to hydrological factors. The model does a good job of simulating the change in the size of the lake. But it does not simulate the vulnerability of the ecosystem to changes in volume. Rather, it relies on elevation to serve as a proxy for the health of the ecosystem. You may deal with this limitation in the "model merger" exercise at the end of the chapter.

Exercises

1. **Verification:**

   Build the fourth model, set DT to 0.25 years, and simulate the model to verify that it matches the results in this chapter. Document your work with a printed copy of the flow diagram and time graphs to match the results in figure 4.13. If you are not sure about all the equations, double-check your equations against selected equations below. For example, here are the equations describing the evaporation and the specific gravity:

```
evaporation = surface_area*evaporation_rate
evaporation_rate =
fresh_water_evaporation_rate*evaporation_rate_multiplier_from_specfic_gravity
specific_gravity = (water_in_lake*mass_of_fresh_water+total_dissolved_solids)/
(water_in_lake*mass_of_fresh_water)
```

Here are the various constants:

```
diversion_inflows = 1.6
evapotranspiration = 13
export = 100
exposed_lake_bottom_evaporation = 12
fresh_water_evaporation_rate = 3.75
ground_water_export = 7.3
mass_of_fresh_water = 1.359
net_grant_lake_evaporation = 1.3
net_land_surface_precipitation = 9
non_sierra_runoff = 20
precipitation_rate = 0.667
sierra_gauged_runoff = 150
total_dissolved_solids = 230
ungauged_sierra_runoff = 17
```

Here are the three graphs for the nonlinear relationships:

```
elevation = GRAPH(water_in_lake)
(0.00, 6224), (1000, 6335), (2000, 6369), (3000, 6392),
(4000, 6412), (5000, 6430), (6000, 6447), (7000, 6463), (8000, 6477)

evaporation_rate_multiplier_from_specfic_gravity = GRAPH(specific_gravity)
(1.00, 1.00), (1.05, 0.963), (1.10, 0.926), (1.15, 0.88), (1.20, 0.833), (1.25, 0.785),
(1.30, 0.737), (1.35, 0.688), (1.40, 0.64)

surface_area = GRAPH(water_in_lake)
(0.00, 0.00), (1000, 24.7), (2000, 35.3), (3000, 48.6), (4000, 54.3),
(5000, 57.2),(6000, 61.6), (7000, 66.0), (8000, 69.8)
```

2. **Check the dynamic equilibrium:**

   Look at the numerical values for each of the flows at the end of the simulation in the previous exercise. Add the flows into and out of the stock. Is the lake in dynamic equilibrium?

3. **Simulate a buffer policy:**

   Modify the model to allow it to simulate a "buffer policy." The idea is to specify a target range for the lake elevation where you are confident that the ecosystem is "safe" as long as the elevation is confined to the range. (A similar policy was advocated by the Mono Lake Committee.) For this exercise, assume a 10-foot buffer zone as shown below.

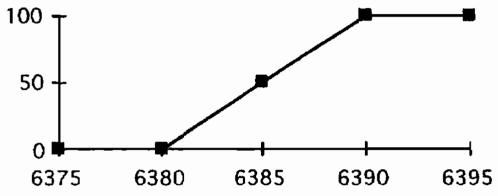

Expand the model to allow exports to depend on the elevation of the lake. If the *elevation* is below 6,380 feet, *export* must be zero. If the *elevation* exceeds 6,390 feet, *export* is set at the "historical" value of 100 KAF/yr. If the elevation is somewhere inside the buffer zone, the *export* should change in a linear manner. Simulate the model over the time period from 1990 to 2010. Document your work with a copy of the flow diagram, the equations, and a graph of *elevation* and *exports*.

4. **Monthly version of the buffer policy model:**

   Change the interpretation of "time" from years to months and repeat the previous test of the buffer policy. Let time run from 0 to 240 months (month 0 corresponds to the year 1990 in the previous simulation). Retain the convention of measuring model inputs in KAF/yr or in feet/yr, but make sure your equations have the correct units. Set the DT to 1 month; simulate the model; and turn in a time graph of ***elevation*** and ***exports*** with the same scales as in the third exercise. You should get the same results as before.

5. **Buffer policy with seasonal changes in runoff:**

   Expand the monthly model from the fourth exercise to allow the *sierra gauged runoff* to vary from one month to another within the year. Use the "monthly counter" (see appendix H) to represent the months of a year. Most of the runoff occurs from melting snowpack during the spring and summer, so set these monthly flows to higher than average. Set the fall and winter flows to lower than average. Don't forget to make sure that the average over all 12 months turns out to be 150 KAF/yr or 12.5 KAF/month. Next, you need to make sure that the model does not export water that may not be flowing down the streams in the fall and winter months. So change the export buffer policy to a fraction of the *sierra gauged runoff.* Document your work with a copy of the flow diagram, the equations, a time graph showing elevation and exports with the same scales as in the previous exercises.

6. **Control Board policy:**

   Hart (1996, 171) describes some of the details of the State Water Resources Control Board's 1994 policy to rebuild the lake toward an elevation of 6,392 feet. He explains that the city is allowed no diversions until the lake reaches 6,377 feet. Then it would be allowed 4.5 KAF/yr until the level reaches 6,390 feet. It would be allowed 16 KAF/yr until the lake reaches 6,391 feet. If the lake reaches higher levels, the city could divert all water in excess of fish flows (which are expected to be 30.8 KAF/yr). Draw an export/elevation chart similar to the chart in the third exercise to approximate the Control Board policy. Then simulate the policy. Document your results with a time graph of elevation and exports. How long does it take for the lake to reach the target elevation? How much water is the city able to export once the lake reaches dynamic equilibrium?

7. **Advanced exercise—model merger from home page:**

   It is often useful to merge two models dealing with different aspects of a system because the combined model may teach us something new. That is, we gain insights that could not be gained from operating the two models separately. After you have learned about the use of "conveyors" in models of animal populations (see figure 14.2), turn to the home page to read about a model of the brine shrimp population of Mono Lake. You are to imagine that the new model has been developed independently of the hydrology model in this chapter. Your job is to merge the two models to provide an internally consistent simulation of size of the brine shrimp population as well as the size of the lake.

   You will discover what many modeling teams discover when working in large organizations: independently developed models do not necessarily fit together just because they deal with the same topic. You may need to make some adjustments in one or both of the models in order to combine them into a holistic picture of the lake and its shrimp. After you succeed in merging the models, use the combined model to examine the Control Board policy tested in the previous exercise. Can you design an export policy based on the observed brine shrimp population rather than the observed elevation? Does your new policy allow the city to export more water from the basin?

Post Script

The focus of this chapter is the Mono Basin and the need for reduced diversions in order to rebuild the lake. But many readers will wonder about the wider implications of reducing the diversions. Where will the extra water come from—

- the Colorado River?
- new reservoirs in Northern California?
- desalinization plants along the ocean?

I believe the most attractive supply of new water to compensate for reduced diversions from the Mono Basin is hidden in the homes and businesses of the City of Los Angeles. The inefficient methods of water use within the city provide an investment opportunity for the water department. Rather than investing in new reservoirs hundreds of miles away, the department could commit its financial resources to help its own customers invest in water efficiency. In my opinion, the saved water could more than compensate for the reduced diversions from the Mono Basin. Moreover, efficiency programs can lead to reductions in the customer's average water bill. Chapter 23 explains how conservation programs can be advantageous to electric utilities, and many of the advantages to an electric utility apply to a water utility as well. Regarding the city's plans, the department has announced that it intends to "meet increased water demands created by growth through a combination of water conservation and recycling programs." The department has invested over $50 million in demand-side management between 1990 and 1995, and it plans to continue demand-side program spending at around $10 million per year between 1995 and 2005 (LADWP 1995, ch. 5).

Further Readings

- Information on Mono Lake is given in several reports by academics (Winkler 1977; NRC 1987; Botkin et al. 1988; Hart 1996), by the City of Los Angeles (LADWP 1987, 1988a,b) and by the Mono Lake Committee (1989, 1997).
- The most extensive single collection of information is the Environmental Impact Report prepared for the State Water Resources Control Board's review of the Mono Basin Water Rights of the City of Los Angeles (Jones and Stokes Associates 1993).
- Much of the modeling information is taken from the water balance model developed by Peter Vorster. The model is thoroughly explained and tested in his thesis (Vorster 1985).
- Supplemental information may be found on this book's home page.