BWeb Copy of the Aluminum Chapter from the 1st Edition of *Modeling the Environment*

Chapter 19. Volatility in Aluminum Production

Aluminum smelting is a commodity industry with highly volatile prices. This chapter uses system dynamics to look "inside the industry" to explain part of the volatility. Aluminum smelting is of special interest to environmental students because of its huge electricity requirement and the potential to reduce electricity consumption through recycling. Aluminum smelting is also of special interest as an example of a commodity production cycle. The chapter begins with an aggregate model of the aluminum industry. It then exploits the power of arrays to develop a model which distinguishes between smelters. It concludes with a discussion of system dynamics and commodity production cycles.

Background

Aluminum is an abundant element comprising about 8% of the earth's crust. (The oxide of aluminum is found in nearly all common rocks.) Yet for most of the nineteenth century, aluminum was a precious metal, costing over $500 per pound (Smith 1988, p. 2). Aluminum was an economic paradox, naturally abundant but extremely expensive. The paradox arose from the technical difficulty in separating aluminum from the oxide of aluminum (alumina).

The technical challenge was overcome in 1886 by Charles Hall and Paul Heroult. Hall was a chemistry student at Oberlin College; Heroult was an inventor at the Ecole de Mines. Their experiments focused on electrolysis to remove aluminum metal from alumina. The challenge was to find a stable bath of molten salts so that electrolytic reduction could proceed on a continuous basis. An electric current was established from a carbon anode inserted into the bath to the carbon lined walls of the pot holding the bath. Alumina was fed into the bath, and metallic aluminum precipitated on the walls of the pot. (The walls were lined with carbon to serve as the cathode in the electrolytic reaction.) Hall and Heroults' experiments provided the basis for the "Hall-Heroult Process" which continues to this day as the only viable means for smelting aluminum on a commercial basis (Smith 1988, p. 17).

In today's industry, a large smelter might produce around 0.2 million metric tons (mmt) of aluminum each year. The smelter would be located close to a cheap source of electric power, typically a large scale hydro-electric facility. With 2,204 pounds in a metric ton, this smelter would produce about 440 million pounds of ingots each year. If ingots were priced at $1 per pound, annual revenues would be 440 million $/yr. The smelters spread around the world have a total capacity exceeding 16 mmt/yr. With ingots at $1 per pound, world wide revenues would exceed 35 billion $/yr.

Aluminum and Electricity

Aluminum smelting is extremely energy intensive, and it is important to understand the close connection between aluminum production and electricity generation. In the early years following Hall's discovery, the Pittsburgh Reduction Company held Hall's patent. They used electricity generated from burning coal and gas to run a small smelting works in New Kensington, Pennsylvania, but they knew that a growing industry would require cheap and massive quantities of electricity. In June of 1893, they built a new smelting works near Niagara Falls (see drawing) and became the first customer of the Niagara Falls Power Company. The partnership between large scale aluminum production and power generation was established. The Pittsburgh Reduction Company succeeded in building both their production capacity and in building a market for the new metal. This small company grew into the Aluminum Company of America, one of the most successful monopolies in American history (Smith 1988).

To appreciate the electricity requirements of aluminum smelting, imagine a modern smelter with 0.2 mmt/yr of capacity and pot lines operating at 7 kilowatt hours (kwh) per pound. This smelter would consume around 3 billion kwh/yr. To place this power requirement in perspective, let's imagine that the electricity comes from a dam with 100 feet of head over the turbines. Each acre-foot (AF) of water passing through these turbines generates 870 kwh of electricity. To keep all the pot lines in operation, the company would need 3.4 million AF/yr passing through the turbines, a huge flow of water (over five times larger than the annual demand for water in the City of Los Angeles). The huge power requirement of smelting makes the aluminum industry an interesting case study for students of the environment. The aluminum in recycled products can be converted to metal ingots with only around 5% of the electricity required in primary production. You may simulate the energy savings from recycling in the exercises at the end of the chapter.

Volatility in Aluminum Prices

Figure 19.1 shows the price of aluminum ingots during the 1970s and 1980s (Knight Ridder 1993, p. 4). The time series shows major swings in the price in a relatively short time interval. Near the end of the 1980s, for example, the ingot price increased to over 100 cents/pound. But within the next two years, the price had fallen to only 50 cents/lb. Another upswing in the late 1980s brought the price to back over 100 cents/pound, but the price returned to around 60 cents/pound by the end of the decade. The large price variations pose difficult problems for the smelter operators as well as their customers. Smelters are complex facilities with a highly trained work force and massive electricity requirements. The smelter owners can't simply turn the smelter on with each upswing in price or turn them off with each downswing. So the high price volatility in Figure 19.1 poses the possibility of highly volatile profits/losses from smelter operation.

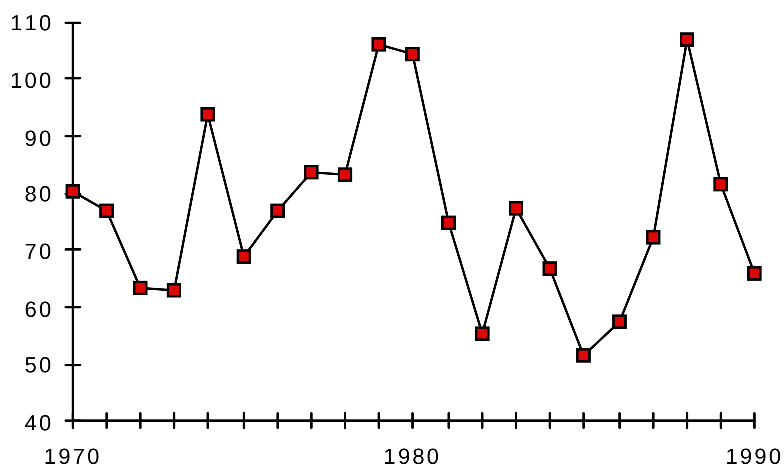

Figure 19.1. Price of aluminum ingots in cents/pound, in constant 1987 dollars.

Figure 19.2 summarizes the costs of operating smelters around the world. This graph is organized with cumulative smelting capacity on the horizontal axis and the variable cost of operation on the vertical axis. This particular curve has adapted from a confidential review of the world industry conducted in 1994. Variable costs are reported in cents/pound, using 1993 US dollars. They are comprised mainly of the cost of alumina, electricity and labor. Smelting capacity is based on the 146 smelters in the "western world." (The review did not include smelters in eastern Europe and former USSR.)

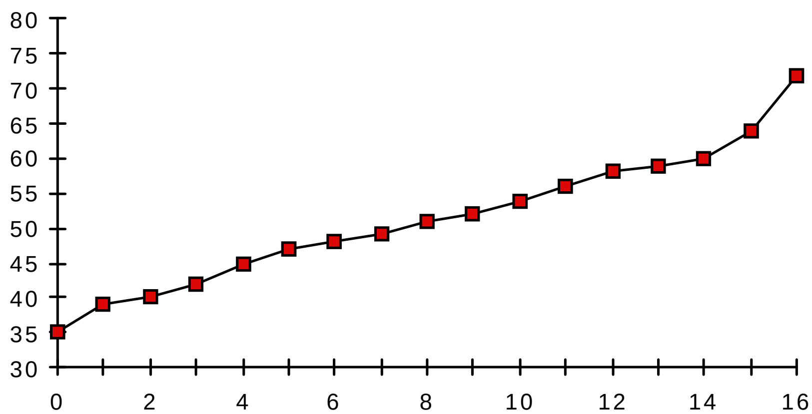

Figure 19.2. Industry cost curve with variable cost (cents/pound) on the Y axis and cumulative smelting capacity (mmt/yr) on the x axis.

This single curve summarizes costs from 146 different smelters arranged to show the gradual increase in variable costs as we proceed from the cheapest to the most expensive smelters. The low cost smelters are located in regions with cheap electricity. Many of these are in the "ABC" countries: Australia, Brazil and Canada (Peck 1988, p. 15). In Australia, the power comes from the recent development of large coal fields. In Brazil and Canada, the power comes from massive hydro-electric development. At the time of the survey, there were 32 smelters in the USA, many located close to large scale hydro-electric development (i.e., in the Tennessee Valley and along the Columbia River.) The highest cost smelters are located in countries whose power rates have increased and which pay their workers higher hourly wages. They tend to be located in western Europe where, despite its high cost, aluminum is supported by public policy and government subsidies (Peck 1988, p. 82).

To interpret the industry cost curve, imagine that you manage a smelter with a variable cost of 60 cents/pound. According to Figure 19.2, there are plenty of smelters in the world with lower costs; their cumulative capacity is 14 mmt/yr. And there are some smelters with higher operating costs; their combined capacity amounts to 2 mmt/yr. Now suppose the market price for ingots were to remain constant at 61 cents/pound. You could operate your smelter and earn 1 cent/pound. You would expect the 14 mmt/yr of lower cost capacity to be in operation as well, and you would probably expect the remaining 2 mmt/yr of capacity to be idle. If the total demand for ingots were 14 mmt/yr, annual demand would be satisfied, and the price might remain at 61 cents/pound in the future.

But Figure 19.1 shows that the price of ingots is not inclined to remain constant over time. Rather, we see major price fluctuations by plus-or-minus 100% in just a few years. If you think about your 60 cent/pound smelter, you will immediately see major opportunities. If your smelter were operating near the end of the 1970s, you could sell ingots at 110 cents/pound, earning around 50 cents/pound after covering your variable costs. But what would you do in the early 1980s when the price dips below 60 cents/pound? Would you close the smelter in 1982? Would you reopen the smelter in 1983? Would you shut it down again in 1985?

These are not abstract questions to set the stage for the modeling example. They are serious questions whose answers impact the livelihood of thousands of families across the country. The problem is particularly acute in the northwest, as indicated by the following headlines:

Northwest aluminum approaches meltdown.
The Oregonian, Feb. 6, 1994
Reynolds rides roller coaster: ex-workers wait.
The Oregonian, April 27, 1995.

This chapter focuses on the underlying causes of the price volatility and its potential to turn the smelting industry into a "roller coaster" industry.

External Versus Internal Sources of Volatility

The volatility in the aluminum industry arises from the interplay of numerous factors, but no single factor is more important that the price of electricity. Some experts attribute the high price volatility in the 1970s to major changes in the price of electricity during the decade of the “energy crisis.” This was a decade in which skyrocketing oil prices triggered major increases in all fossil fuel prices and subsequent increases in the costs of electricity. Smelters in areas served by higher cost power stations experienced major increases in operating costs. The smelters in Japan were particularly vulnerable, and Japan closed down 75% of their smelting industry during the 1980s (Peck 1988, p. 90). With a combination of permanent and temporary closures, is it any wonder that aluminum prices increased to such high levels in the 1970s?

Suppose we attribute the volatility in the 1970s to the “energy crisis,” an external factor beyond the normal range of managerial control of the aluminum industry. But what about the volatility in the 1980s? Is there another external event of similar importance to the “energy crisis”? Many experts argue that the break down of the Soviet Union is the key event of the 1980s. The dissolution of the USSR caused a massive reduction in the soviet military which consumed a major share of the aluminum produced in Russia. The Russian smelters turned to the western world, and many experts attribute the low aluminum prices in the 1980s to the extra production entering the western market.

These arguments looks to factors external to the industry. Whether it’s the “energy crisis” or the breakdown of the Soviet Union, we are looking at factors beyond the normal range of control of the industry leaders. Of course, the volatility in any industry is due to a combination of external forces and the internal workings of the system. It may help to think of other systems that are subjected to highly variable, outside forces. Recall the discussion of homeostasis from Chapter 8, and consider the example of temperature control in the human body. The human body is frequently subjected to outside disturbances, but body temperature is normally maintained at around 98.6 degrees F. To understand why this is possible, we need to look “inside the system” to understand the feedback mechanisms (as in Appendix K). In contrast to the human body, the world aluminum industry seems unable to withstand the impact of external events without major swings in market prices and the ensuing “roller coaster” impacts on workers and their families.

An Initial Model

Figure 19.3 shows an equilibrium diagram for an initial model. One stock is used to keep track of the aluminum held in inventory at the smelters and the mills around the world. A second stock represents the aluminum products in use. The model distinguishes between primary production at smelters and secondary production from the recycling of used products. Stocks are measured in millions of metric tons (mmt) of metallic aluminum. Time is measured in months, and each of the flows is measured in mmt/month. Annual amounts are reported to aid in interpretation. For example, *primary production* is 14 mmt/yr; *secondary production* is 2 mmt/yr and *annual demand* is 16 mmt/yr. Starting at the top of Figure 19.3, the world smelting capacity is assumed to be 16 mmt/yr. (The model does not include the smelters in the former USSR). The fraction of smelting capacity in operation depends on the producer’s view of the ingot price. Figure 19.3 shows both the *ingot price* and the *producers lagged price* at 60 cents/pound.

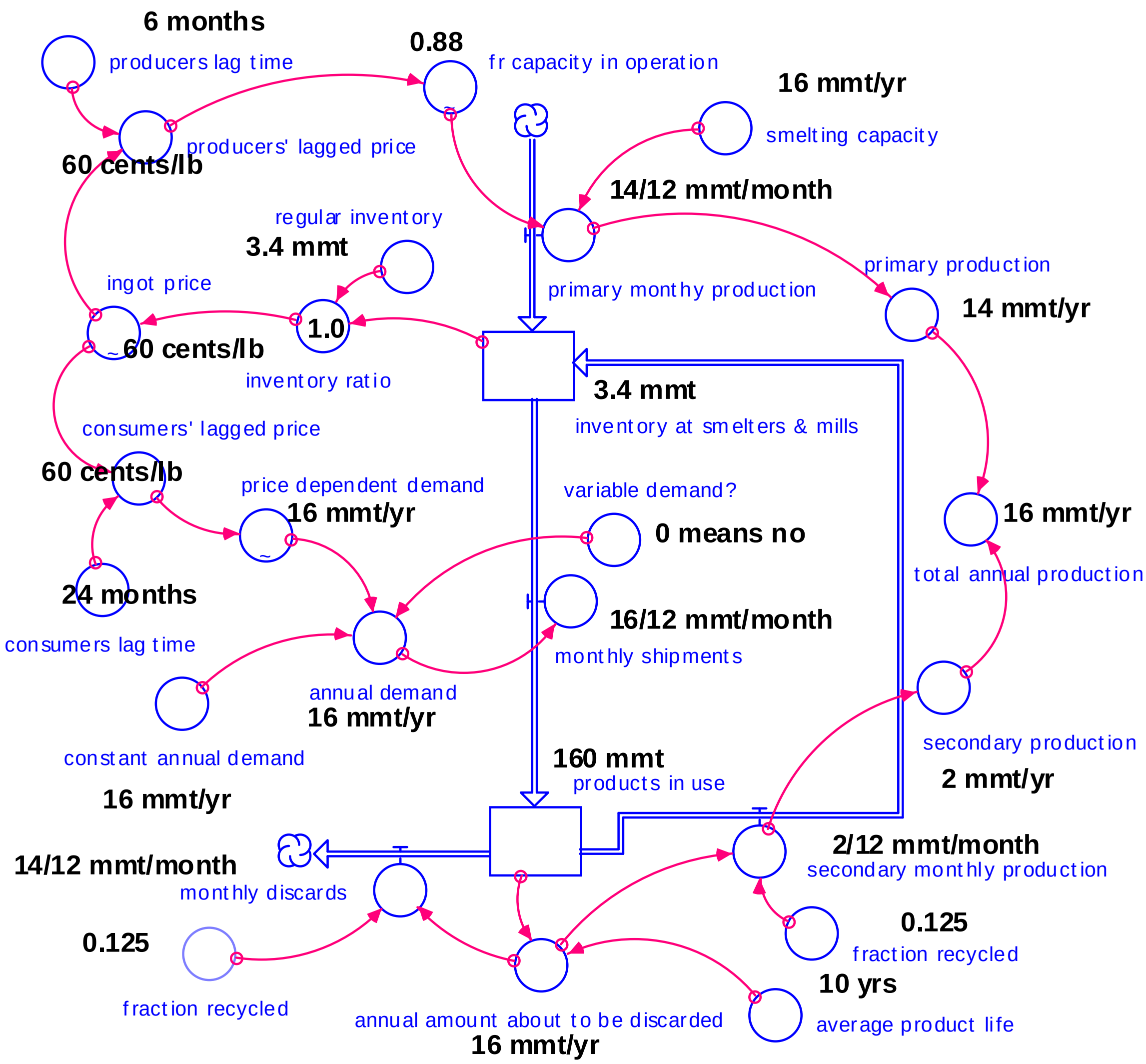

Figure 19.3. Aluminum model with equilibrium conditions.

The fraction of smelting capacity that would operate at 60 cents/pound is found by a nonlinear graph (~) shaped to match the industry cost curve in Figure 19.2. At 60 cents/pound, for example, 88% of the world's smelting capacity would be in operation. Primary production would be 14 mmt/yr, building the stock of inventory held by the smelters and the mills around the world. The inventory is initialized at 3.4 mmt, and this amount is taken as the *regular inventory* needed to ensure efficient operation (Metal Stats 1996, p. 123). You can see from Figure 19.3 that the inventory would be maintained at 3.4 mmt because there is 14 mmt/yr of *primary production* and 2 mmt/yr of *secondary production*, exactly the amount needed to satisfy the *annual demand* of 16 mmt/yr. The annual demand is fixed at 16 mmt/yr by setting the*variable demand?* to zero. The ? reminds us that this is a binary variable, as explained in Appendix H.

The *monthly shipments* move the aluminum from the mill to the ultimate consumers. The model combines the many uses of aluminum (i.e., automobiles, airplanes, housing, beverages, etc.) into a single category with an *average product life* of 10 years. The total stock of aluminum *products in use* is initialized at 160 mmt. With a ten year product life, around 16 mmt/yr would be ready for discard in the first year. The user specifies the fraction of the potential discards that will be collected and shipped to secondary producers. The US has a relatively well developed system for collection, and secondary production amounts to about 16% of total demand (Metal Stats 1996, p. 318). The *fraction recycled* in Figure 19.3 is set at 12.5% on the assumption that recycling is less developed in world as a whole. The recycled products are delivered to secondary producers who extract the aluminum and sell the metal to the mill operators.

You should be able to write the equations for most of the variables in Figure 19.3 based on the names and the numbers in the equilibrium diagram. The model includes three nonlinear relationships represented by graph functions. For example:

```
fr_capacity_in_operation = GRAPH(producers'_lagged_price)
(20.0, 0.00), (30.0, 0.00), (40.0, 0.19), (50.0, 0.63), (60.0, 0.88),
(70.0, 0.97), (80.0, 1.00)
```

which shows the fraction of smelting capacity that is expected to be in operation based on the *producers' lagged price*. The graph is designed to match the industry cost curve shown in Figure 19.2. If producers base their decision on 80 cents/pound, for example 100% of the smelters would be in operation. If the producers' price were to fall to 30 cents/pound, none of the smelters would operate. The next nonlinear relationship controls the demand for ingots:

```
price_dependent_demand = GRAPH(consumers'_lagged_price)
(20.0, 24.0), (30.0, 22.0), (40.0, 20.0), (50.0, 18.0), (60.0, 16.0),
(70.0, 14.0), (80.0, 12.0), (90.0, 10.0), (100, 10.0), (110, 10.0), (120, 10.0)
```

The demand for aluminum may change with changes in the consumers view of price. If the price were to increase from 60 to 70 cents/pound, for example, this graph would lower the demand from 16 to 14 mmt/yr. This example assumes a 10/60 or 17% increase in price would trigger a 2/16 or a 13% reduction in demand. If we compare the relative changes, we would say that the "price elasticity of demand" is -0.75. The third graph represents changes in the *ingot price* with changes in the inventory ratio:

```
ingot_price = GRAPH(inventory_ratio)
(0.00, 200), (0.25, 190), (0.5, 150), (0.75, 100), (1.00, 60.0),
(1.25, 50.0), (1.50, 45.0), (1.75, 40.0), (2.00, 40.0)
```

A ratio of 1.0 means that the industry has approximately the inventory needed to allow efficient operations of the smelters and the mills. The model assumes that ingot prices would be at 60 cents/pound under these "regular" conditions. Lower values of the inventory ratio are assumed to push the price higher. If the ratio falls to 75%, for example, ingot prices are assumed to increase to 100 cents/pound. If inventories fall to 50% of the regular levels, ingot price is assumed to climb even higher to 150 cents/pound. The shape of this graph function is taken generally descriptive of inventory and price changes reported over 1985 to 1995 in a survey of inventories maintained at the London Metals Exchange (Industry Surveys 1995).

The model uses two lagged relationships. The first involves the producers reaction to a change in prices. Let's assume that smelter operators watch price changes over a time interval before committing to opening or closing a smelter. Closing costs could be several million dollars (Peck 1988, p. 10), so we should assume that operators will not react to each and every fluctuation in the ingot price. It makes more sense to assume that they react to a time averaged price which can be represented by a third order smooth function. The length of the lag is uncertain, but let's begin with a value of 6 months.

```
producers'_lagged_price = smth3(ingot_price,producers_lag_time,60)
producers_lag_time = 6
```

Consumers are not likely to react instantaneously to price changes either. The initial model assumes that the delay in the consumers' reactions may be described by a third order smoothing delay with a two year lag time:

```
consumers'_lagged_price = smth3(ingot_price,consumers_lag_time,60)
consumers_lag_time = 24
```

The third entry in the SMTH3 function is 60 for both the producers and the consumers. This means that the simulation begins with both producers and consumers using 60 cents/pound as the appropriate price for decision making.

Simulating a Production Cycle

The equilibrium conditions will persist from one year to another if you simulate the model over time. Every year will show a demand for 16 mmt/yr of aluminum. Total production will exactly balance the demand; ingot prices will remain constant at 60 cents/pound, and 88% of the world's smelting capacity will remain in operation. To learn if this equilibrium is stable, we should introduce a disturbance. One way to disturb the equilibrium is to assume that the industry begins the simulation with extra inventory.

Figure 19.4 shows the simulation results with the initial inventory set at 4.4 mmt rather than 3.4 mmt. The *annual demand* for aluminum is held constant at 16 mmt/yr to allow us to concentrate on a situation in which only the producers react to changes in prices. In other words, the initial simulation is designed to test for a "production cycle" in aluminum. Figure 19.4 shows that producers begin the simulation producing 16 mmt/yr, exactly the amount needed to balance demand. The simulation runs for 96 months to allow sufficient time to see if there are volatile swings in prices. This simulation reveals a production cycle which may be attributed entirely to the operating decisions of the primary producers. *Total demand* and *secondary production* are both constant throughout the simulation. *Primary production* is simulated to decline in the first few months of the simulation because of the low ingot prices. The decline in *primary production* causes total production to fall below demand during the first few months. This allows the inventory to decline to more regular levels, and prices are simulated to swing upward. Producers react to the upswing after a delay, and more capacity comes into operation.

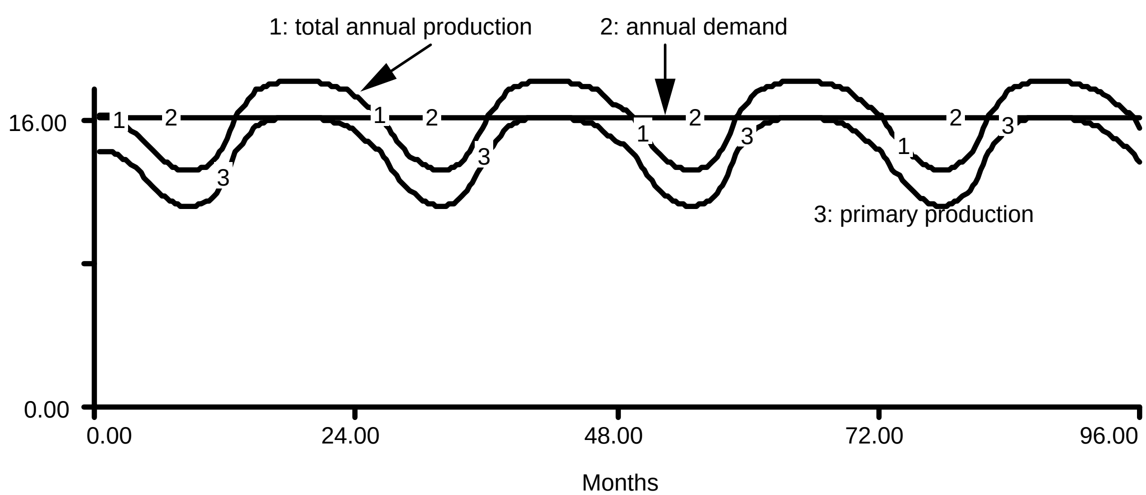

Figure 19.4. Initial simulation of aluminum production and demand.

Figure 19.5 shows that almost all of the smelting capacity would be in operation shortly after the peak in ingot prices. This high production causes inventories to build past the regular levels causing the drop in prices seen around the second year of the simulation. Smelters are gradually taken out of operation during this time period, and the excess inventories are reduced. Figure 19.5 shows a second upswing in prices around the 30th month of the simulation, and the smelters' response is similar to their response to the first upswing in prices.

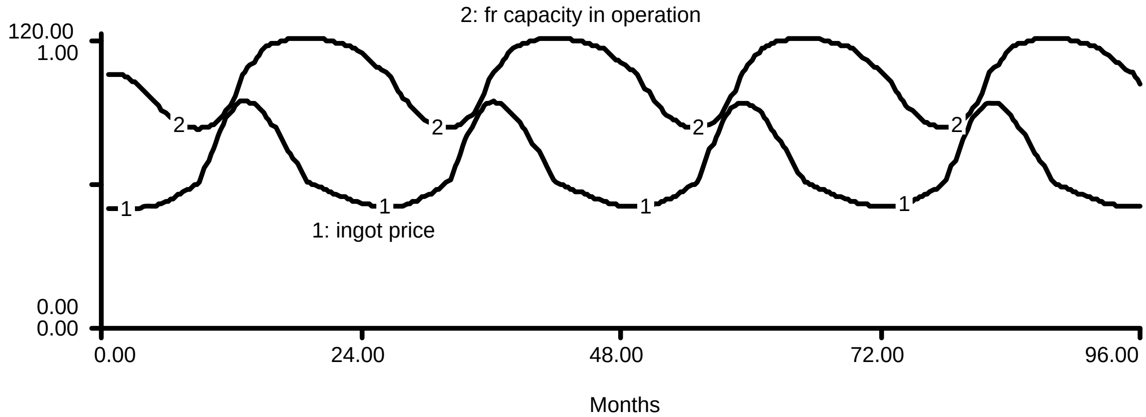

Figure 19.5. Initial simulation of ingot prices and capacity operation.

This simulation reveals that the action of primary aluminum producers could be a major contributor to the volatility in the industry. The simulation shows price swings from a low of around 50 cents/pound to a high of over 90 cents/pound. These large swings are quite surprising when you consider that there are no further disturbances after starting out the simulation with extra inventory. Moreover, the price swings occur even though there are no variations in the demand for aluminum, in the amount of secondary production, or in the total smelting capacity.

Figure 19.6 shows the negative feedback loop involving the reaction of primary producers to changes in aluminum price. Primary production adds to the inventory at smelters and mills. This builds the inventory ratio and lowers the ingot price. After a delay to watch and evaluate the price changes, the primary producers reduce the fraction of capacity in operation. This lowers monthly production and allows monthly shipments to lower inventory levels to regular levels. The // on the link from the ingot price to the fraction of capacity operation draws our attention to the key delay in the loop. You learned in Chapter 17 that the length of the time lag can be a crucial factor controlling the stability of an oscillatory system. So it is important to test the sensitivity of the model to changes in the length of the producers lag time.

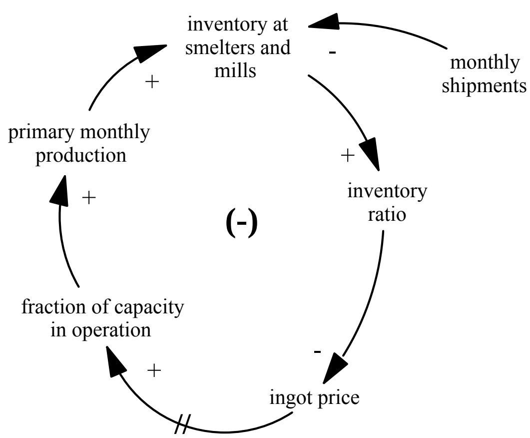

Figure 19.6. Negative feedback responsible for the simulated production cycle.

Figure 19.7 shows three simulations of the aluminum production cycle.
The first run assumes that producers average the ingot prices for 3 months before opening or closing their smelters. With this lag time, the industry is simulated to react in a highly stable manner. The second run in

Figure 19.7 is the same as the simulation shown previously. It shows that a 6 month lag would create sustained oscillations with a period of around two years.

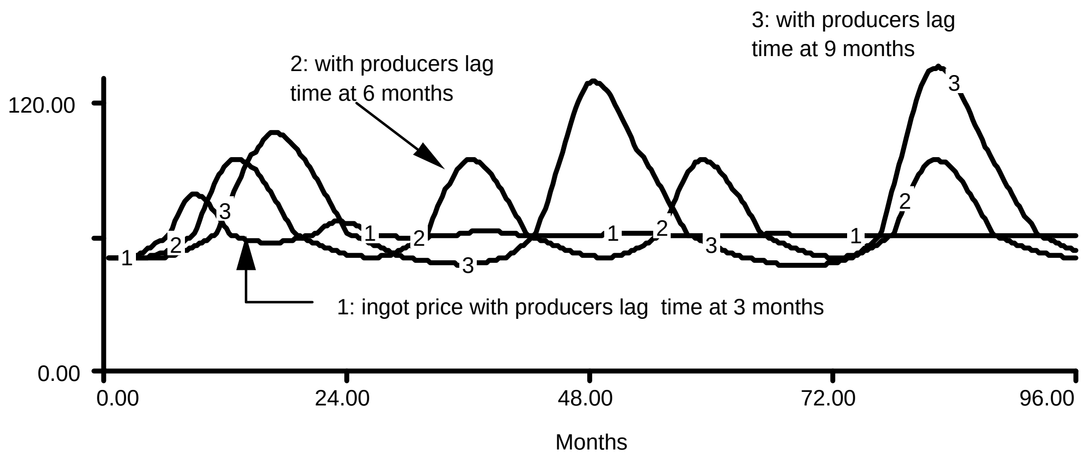

Figure 19.7. Simulated production cycle with three values of the lag in producers' reaction to ingot prices.

The third run assumes that producers average the price over 9 months before committing to opening or closing of their smelters. This longer lag time causes more volatile cycles with a period of around three years. Figure 19.7 demonstrates that the length of the producers' lag time is an important determinant of the overall stability of the aluminum industry.

Adding the Consumers' Reaction

Figure 19.8 expands the causal diagram to show the reaction of consumers to changes in ingot prices. An increase in price will cause a reduction in consumption and a corresponding reduction in monthly shipments. When less aluminum is shipped, the inventory tends to build more rapidly causing a reduction in the price.

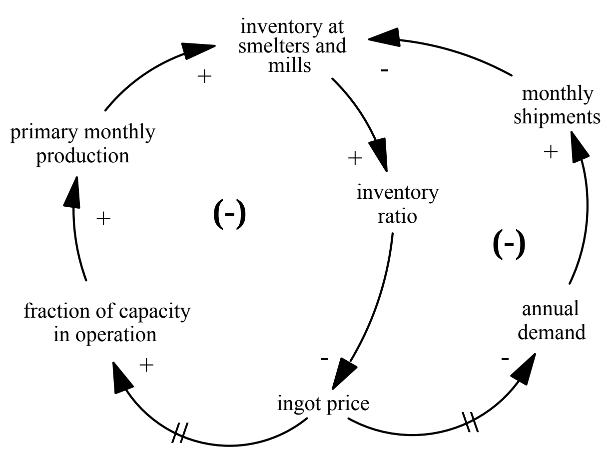

Figure 19.8. Coupled negative feedback loops controlling the supply and demand for aluminum.

The demand loop may be activated by changing the *variable demand?* input from 0 to 1. With this change, the *annual demand* will be based on *price dependent demand* by the following equations:

```
variable_demand? = 1
annual_demand = if (variable_demand?=1) then price_dependent_demand
else constant_annual_demand
```

Recall that the price induced change in demand corresponds to a price elasticity of -0.75. Long run price elasticities have been estimated to range from -0.78 (in construction) to -1.60 (in transportation) by Charles River Associates (CRA 1971). (Their study was completed during a period prior to the high volatility in prices shown in Figure 19.1.) The CRA study also sheds light on the time interval required for the full price response to appear. CRA used an econometric model with a Kyock (1954) lag structure to distinguish between the short run and long run price elasticities. With this approach, the ratio of the two elasticities provides an indication of the length of the consumers' lag time. Their study confirmed what you might expect -- the time lag for consumers to react to changes in ingot prices is much longer than the lag time for producers. Let's set the lag to 24 months. To put this assumption in perspective, imagine that the price of ingots were to increase by 10% today. Consumers are assumed to react by lowering aluminum consumption by 7.5%, but you will have to wait 24 months before seeing the roughly half of their reaction.

Figure 19.9 shows the simulated changes in production and consumption when the consumers' price response is added to the model. The *annual demand* is no longer constant at 16 mmt/yr. Instead, we see modest fluctuations in demand as consumers react to the changes in ingot prices. *Total annual production* varies in a cyclical fashion, due to cycles in *primary production*.

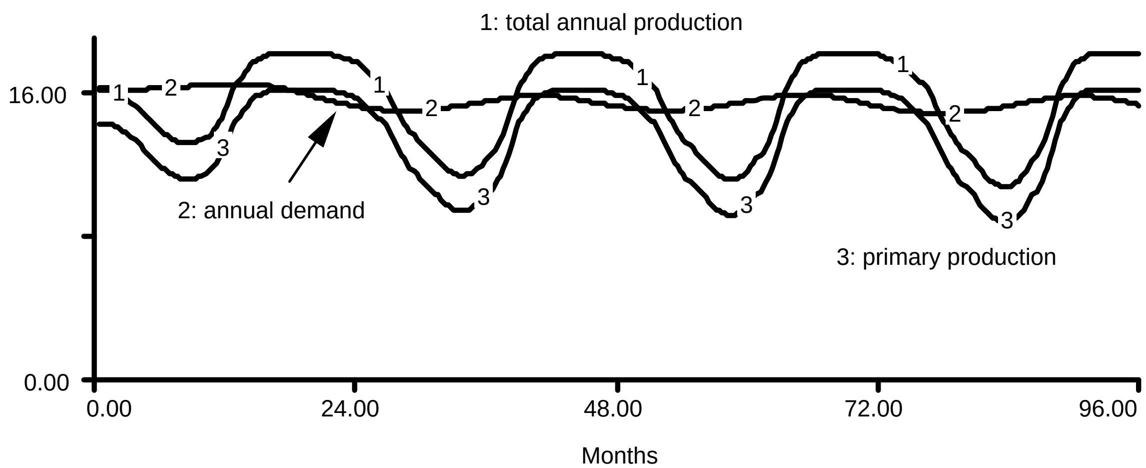

Figure 19.9. Production and demand in a simulation with both producers and consumers reacting to ingot prices.

Figure 19.10 compares the simulated changes in ingot prices in the two simulations. The comparison shows that the introduction of consumer response to prices introduces somewhat more volatility to the system. With variable demand, the cycles show somewhat higher peaks and a somewhat longer period.

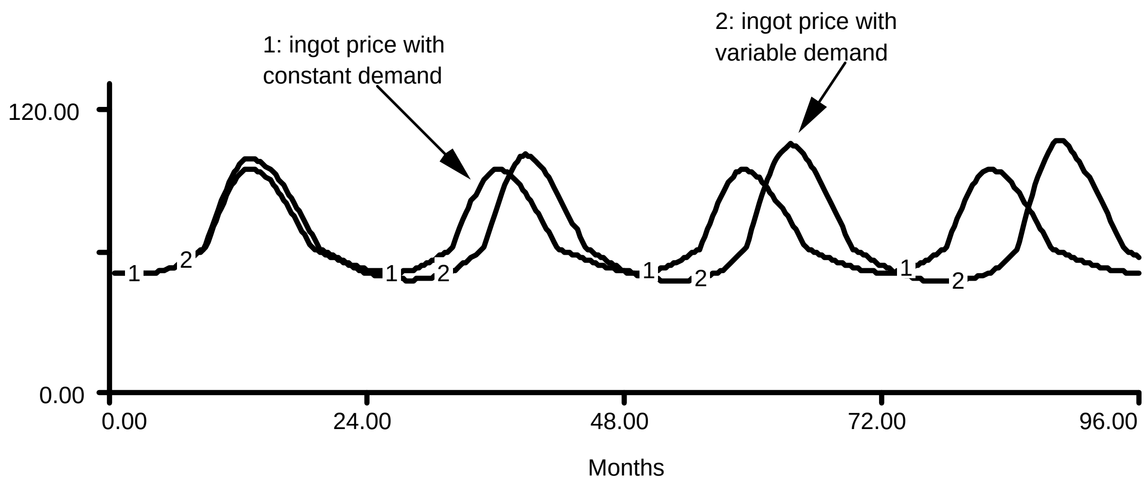

Figure 19.10. Comparison of ingot prices in previous two simulations.

The simulation with variable demand is the more realistic of the two simulations. It shows cyclical behavior in ingot prices with the cycles taking on the characteristic look of a limit cycle. You have read about limit cycles in Chapter 1, but this is the first opportunity to see their shape over time. You can spot limit cycles by their clipped or nonlinear appearance. Another way to check for a limit cycle is through a scatter graph designed to portray the physical limit, as shown in Figure 19.11.

Figure 19.11 displays the same simulation shown in Figure 19.9, but this diagram shows the ingot price versus the fraction of capacity in operation from each month of the simulation.   The starting point is 88% operation and 50 cents/pound. The first few months show a drop in the fraction of capacity in operation. By the time the fraction reaches 70%, the ingot price is increasing, so the dots in Figure 19.11 “change course” from a southerly to an easterly direction. The dots then circle up toward 100% and come back around. By the time they have completed one cycle, the system is further removed from the starting point. This outward growing spiral is characteristic of an unstable system.

Unstable systems can not grow forever; they will eventually encounter limits. In this example, the limits are reached after only one or two cycles. The most visible limit is on the top of Figure 19.11. The fraction of capacity in operation can not exceed 100%. Figure 19.11 also reveals a nonlinear limit on the left side of the diagram. This limitation is somewhat arbitrary. It corresponds to the assumption that aluminum prices are not likely to fall below 40 cents/pound even if inventories build to very high levels.

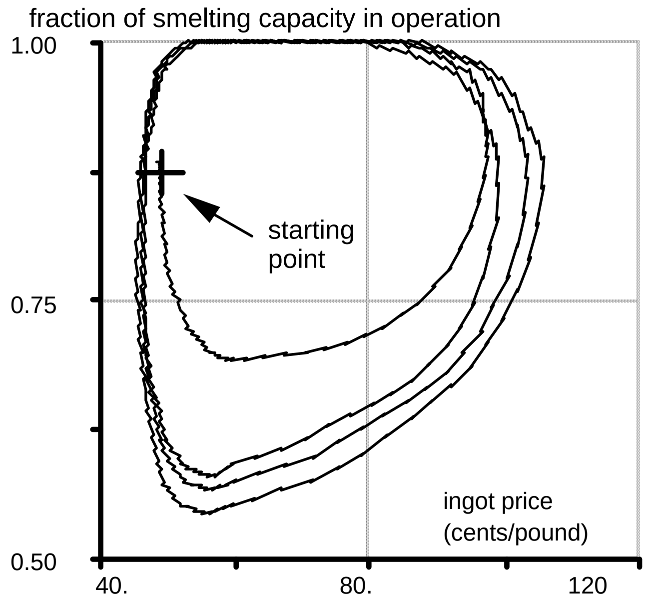

Figure 19.11. Scatter graph depicting the limit cycle in the simulated aluminum industry.

Expanding the Model

The aluminum model can be expanded to provide more realism in several directions (as illustrated in the many exercises). But to continue our focus on price volatility, it seems that we should continue to focus on the operational decision making by primary producers. One way to add realism in this direction is to think about the decision making at individual smelters. There are 146 smelters in the "western world" with widely different operating costs. The previous model combines all of these smelters into one group and focuses on the fraction of the smelters that would be in operation. To bring the new model "closer" to the producers, it would be useful to allow for a smelter-by-smelter simulation. Figure 19.12 shows how this might be done by invoking Stella's array feature for the *variable cost* of individual smelters. This diagram concentrates on the top portion of the flow diagram shown previously.

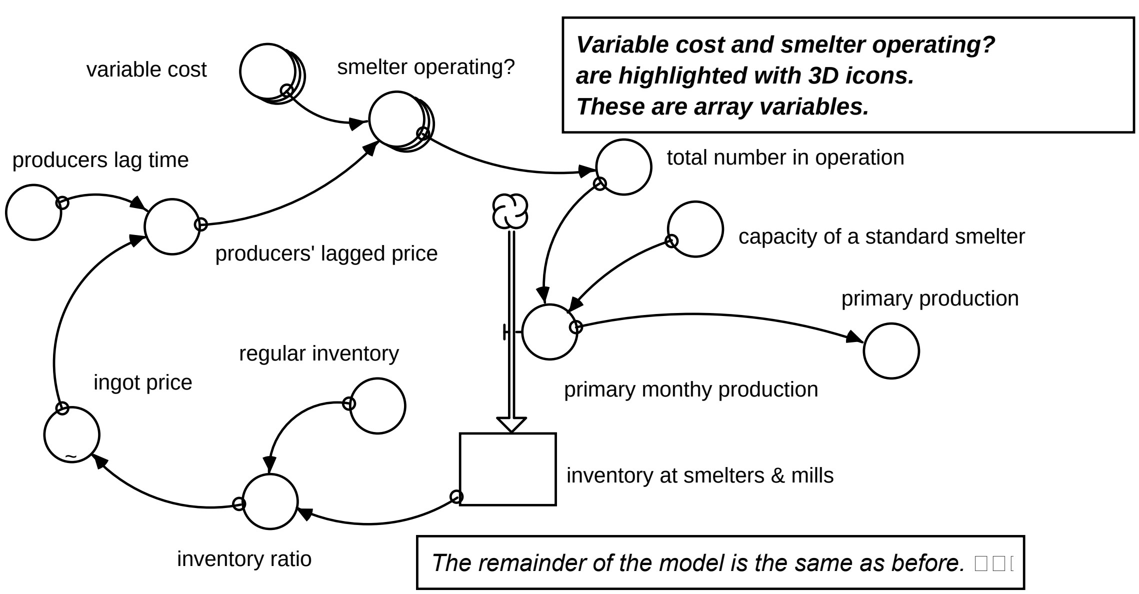

Figure 19.12. New model with array variables for the variable cost and the operating status of individual smelters.

Notice the three dimensional icons for *smelter operating?* and for *variable cost*. The 3-D effect reminds us that these variables are defined as arrays. The model defines an array dimension called *smelter*. Rather than simulating each of the 146 smelters in the "western world," *smelter* will be allowed to run from 1 to 16, as if there are 16 large smelters in the world. (This will be sufficient to demonstrate how arrays are used and to learn if the model yields new insights on price volatility.) Each smelter is assigned a capacity of 1 mmt/yr and is assumed to respond to prices with a 6 month lag. The 16 values of the *variable cost* are defined to match the cost curve in Figure 19.2. The new model compares the lagged price with the variable cost of each smelter to determine whether a smelter is in operation. The ? in *smelter operating?* reminds us that this is a binary variable. Let's define "1" to mean that the smelter is operating. We may then use Stella's ARRAYSUM feature to add the 16 values of the array to learn the total number of smelters in operation. The equations would be written as follows:

```
variable_cost[1] = 35
variable_cost[2] = 40
variable_cost[3] = 42
variable_cost[4] = 45
variable_cost[5] = 46
variable_cost[6] = 48
variable_cost[7] = 49
variable_cost[8] = 50
variable_cost[9] = 52
variable_cost[10] = 54
variable_cost[11] = 55
variable_cost[12] = 56
variable_cost[13] = 58
variable_cost[14] = 60
variable_cost[15] = 65
variable_cost[16] = 70
capacity_of_a_standard_smelter = 1
smelter_operating?[Smelter] =       if(producers'_lagged_price>variable_cost[Smelter]) then 1
else 0
total_number_in_operation = ARRAYSUM(smelter_operating?[*] )
```

Figure 19.13 shows the simulated behavior of the new model. The *annual demand* is constant at 16 mmt/year, so any variability in the *ingot price* may be attributed to the actions of the sixteen smelters. You may compare the new results with the results shown in Figure 19.5 to learn if the new model shows greater volatility in prices. Figure 19.13 shows that ingot prices follow a limit cycle with approximately the same period as in Figure 19.5. The prices at the peak of each cycle are around 120 cents/pound, somewhat higher than the corresponding values shown in Figure 19.5. The greater volatility in the new model is probably caused by the "lumpiness" associated with a simulated industry comprised of 16 large smelters. Verifying whether lumpiness creates increased volatility is left for you as an exercise at the end of the chapter.

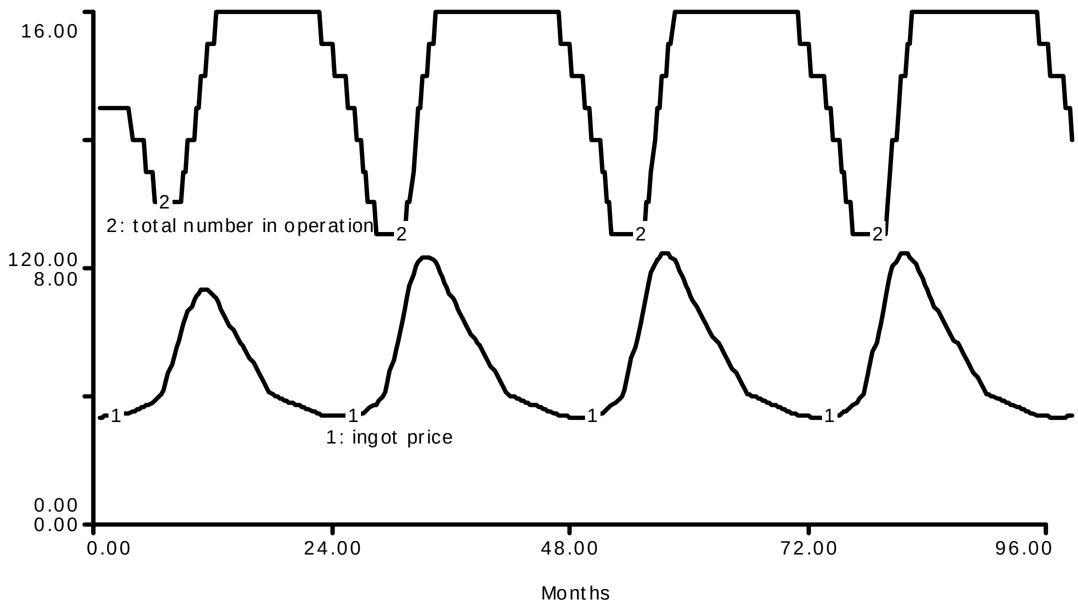

Figure 19.14. Simulated production cycle with 16 large smelters.

Conclusions

This chapter illustrates the use of system dynamics to look "inside a system" for an explanation of highly volatile behavior. The aluminum model indicates that the operating decisions of primary producers could cause substantial price volatility, especially if there are significant lags in their reaction to price changes. If we are to obtain a deeper understanding of industry volatility, it appears that we should concentrate further model improvements in the area of smelter operations. Several such improvements are left as exercises at the end of the chapter.

This chapter focuses on aluminum, but you should know that the approach could be applied to other commodities such as cotton, cocoa, coffee etc. Commodity production cycles are particularly important in developing countries where the export of primary commodities may account for the bulk of the foreign exchange. Meadows (1970) has published a model to explain cyclical production in a hypothetical commodity. The cycles arise, in part, from the delays that agricultural businesses face in expanding production capacity. Meadows replaces hypothetical parameters with parameters drawn from the US hog industry, and the model generates a four year cycle resembling the cycle in US hog production. When the parameters are changed from hog farming to chicken farming, the simulations produce a 30 month cycle similar to the cycles in the poultry industry. When the parameters are changed to represent the longer delays in the beef industry, the simulations reveal a 15 year cycle characteristic of milk cows and cattle. His approach is sensitive to the many differences between the industries, but it succeeds in revealing the underlying factors that generate volatile cycles in all three commodities. Meadows suggests that his approach could be extended to other commodities. Arquitt (1995) follows this suggestion by developing a simulation model of commercial shrimp farming in Thailand. The aluminum model in this chapter may also be viewed as an extension of Meadows' ideas.

Exercises

1. Build and Verify
Build the model shown in Figure 19.3 and verify the simulation results shown in Figure 19.4 and Figure 19.9.

2. Sensitivity to Consumers' Lag Time
Conduct a sensitivity test with the consumers' lag time at 12, 18 and 24 months. Document your results with a comparative graph like the one in Figure 19.7. Does a faster consumer response make the system more stable?

3. Build and Verify the Model with Arrays
Build the model in Figure 19.12 and verify the results in Figure 19.13.

4. More Detail on Operating Costs
Expand the model in Figure 19.12 to include the additional arrays shown below. They provide greater detail on the *total variable costs* at each smelter. Assume that all 16 smelters pay 10 cents/pound for labor, 15 cents/pound (of aluminum) for alumina and 5 cents/pound to transport aluminum to market. Assume that each smelter needs 7 kwh of electricity per pound of aluminum. To create some variability in operating costs, let the *cost of electricity to smelters* vary from a low of 0.7 cents/kwh for the first smelter to a high of 5.7 cents/kwh for the 16th smelter. Run the new model and compare its performance with the results in Figure 19.13.

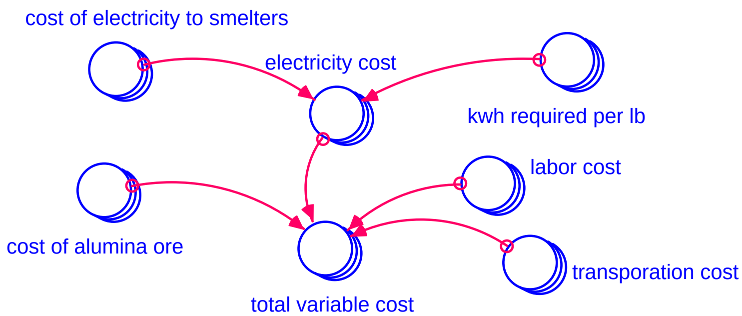

5. More Detail on the Producers' Lag Times
Expand the model in Figure 19.12 to allow for a different lag time at each of the 16 smelters. You may assume that the 6 month lag time describes the decision making at smelters 1-8, the smelters that are likely to remain in operation for most of the simulation. Assume that the remaining smelters are more familiar with shutting down and reopening so they respond to price changes more rapidly. Set their lag time to 3 months, and run the new model to see if you get results similar to Figure 19.14. Which lag time is more important--the 6 month lag used by smelters 1-8 or the 3 month lag time used by smelters 9-16?

6. Keeping One Pot Line Open
Expand the model in Figure 19.12 to allow each smelter to have a different number of pot lines. (The number might range from 4 to 12.) Change the portrayal of the capacity operations to incorporate a requirement to keep one pot line in operation. For example, an unprofitable smelter with ten pot lines will close down nine out of ten pot lines, so only 10% of the capacity will be in operation. Run the new model to learn if this "pot line constraint" changes the simulated volatility in prices.

7. Shut Down Costs
Expand the model in Figure 19.12 to allow the user to specify the shut down costs (in millions of $) at each smelter. Compare the variable cost and ingot prices to determine the monthly losses (if any) at each smelter. Alter the model to simulate closures based on the monthly losses.. Assume that smelters with

higher shut down costs will be willing to absorb longer periods of losses before shutting down their smelters. Run the new model to see if it delivers different insights on the volatility of aluminum prices.

8. Endogenous Price of Alumina

Expand the model in Figure 19.12 to allow the *cost of aluminum ore* to vary with the *ingot price* as shown in the diagram and equations below. Recall that the previous model sets the cost of alumina at 15 cents per pound of aluminum that will be extracted from the alumina. Let's assume that the ratio of 15 cents/pound for alumina relative to a 60 cent/pound price of metallic aluminum is typical of the *metal to ore price ratio*. Define a lagged value of the ingot price with a 3 month lag time to represent short delays in the market adjustments in the cost of alumina.

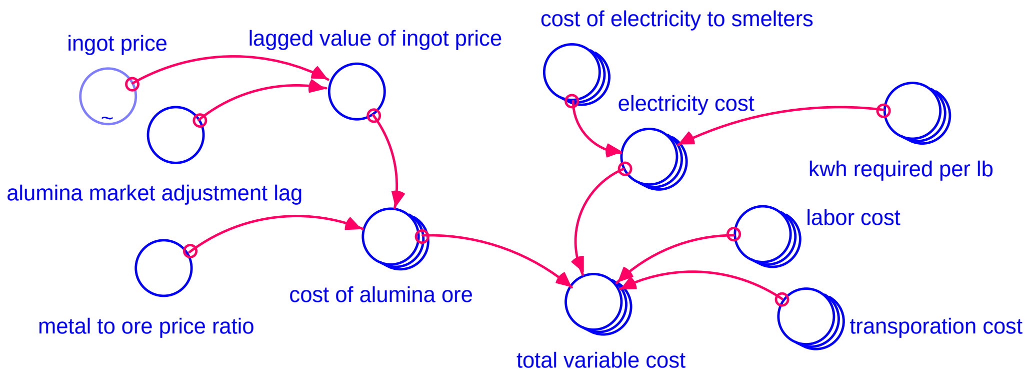

cost_of_alumina_ore[Smelter] = lagged_value_of_ingot_price*metal_to_ore_price_ratio
□lagged_value_of_ingot_price = smth3(ingot_price,alumina_market_adjustment_lag,60)
alumina_market_adjustment_lag = 3
metal_to_ore_price_ratio = .25

Run the new model with different values of the adjustment lag time to learn if the endogenous treatment of alumina costs makes the system more stable.

9. Electricity Consumption

Expand the model from the 4th exercise to keep track of the electricity consumption by smelters. There are 2.2046 pounds in a kilogram and a metric ton is 1,000 kilograms. The model should track of kwh requirements per month and the cumulative kwh for the entire simulation.

Then expand the model to keep track of the electricity consumption by secondary producers which extract aluminum from recycled products. Assume that each pound of secondary aluminum requires only 5% of the electricity needed for primary production (Berk 1982, p. 47). Keep track of the total electricity consumption per month and the cumulative kwh consumed over the entire simulation.

With just a pencil and paper (and perhaps a calculator), estimate the reduction in electricity consumption if we could double the amount of recycling in the model. Record your expected savings in billions of kwh/month and total savings over the entire simulation.

Then change the fraction recycled from 12.5% to 25% at the start of the simulation and run the model to obtain a "systems perspective" on the amount of electricity to be saved by increased recycling. How does the simulated savings differ from the "pencil and paper" estimate? Why are the two estimates different?

10. Causal Loop Diagrams

Expand the causal loop diagram in Figure 19.8 to show any feedback loops created by the recycling of aluminum. Then expand the diagram to show any additional feedback loops if the recycling fraction were to increase with higher ingot prices.

12. Variable Recycling

Expand the model in Figure 19.3 to allow the *fraction recycled* to vary with the ingot price. Rather than 12.5% in every year, let the *fraction* vary from a low of 8%/year to a high of 16%/yr depending on the ingot price. Assume there is no delay in the recyclers' reaction to the price. Does this addition make the system more stable?

Then expand the model from to introduce a 3 month lag in the reaction of the secondary producers to changes in the ingot price. Does the addition of the time lag make the system more stable?

13. End Use Detail

Expand the model in Figure 19.3 with an array to distinguish between different uses of aluminum. Define a new dimension *use* to take on the values:

auto:	for aluminum used in the automobile industry
air	for aluminum used in airplanes
house	for aluminum used in housing construction,
bev	for aluminum used in beverage containers, and
other	for all other uses of aluminum.

Each use should be assigned a different level of demand, and the total demand should match the 16 mmt/yr in the original model. Assign a different price elasticity, a different product lifetime and a different recycling fraction to each use. Set the recycling fractions to ensure that the total recycling is similar to the 2 mmt/yr of secondary production in the original model. Run the new model with constant recycling fractions. Does the detailed treatment of end uses change the volatility of the simulation?

14. New Scrap as well as Old Scrap

The recycling of used products is sometimes called "old scrap." Another source of aluminum is the recycling of "new scrap" which is found on the floors of the mills and fabricators that change aluminum ingots into finished products. In the US industry, new scrap can be twice as large as old scrap (Berk 1982, 51). Expand the model in Figure 19.3 to include aluminum fabrication. Assign one stock to inventories held at smelters and a second stock for inventories held by fabricators (as in Figure 2.10). Split the previous inventory two thirds to the smelters and one third to the fabricators. There is a 26% fabrication loss factor, so, the smelters must deliver 126 pounds to the fabricator for every 100 pounds of finished aluminum products. The extra 26 pounds will fall to the floor as "new scrap." Define a new scrap recycling fraction at 80%, so 80% of the new scrap would be returned to the smelter's inventory. Run the new model to learn if the recycling of new scrap makes the system more stable.

15. Technological Advance

Change the model in the 4th exercise to simulate the investment in advanced technology at the smelters with highly variable operations. Suppose a new technology becomes available in the 48th month of the simulation which reduces the electricity requirement from 7 to 5 kwh per pound and the labor costs by around 4 cents/pound. (Technological possibilities are described by Russell (1981), Jarrett (1987) and INEL (1987). The illustrative benefits in this exercise might be obtained if we were to perfect the "stable cathode/inert anode process".) Run the new model with the assumption that smelters 12-16 adopt the new technology when it becomes available. Does their investment allow them to operate their smelters on a more constant basis?

Then repeat the analysis with the assumption that all 16 smelters invest in the new technology when it becomes available. Does the technological advance allow smelters 12-16 to operate on a more constant basis? Do you see any changes in the fraction of the time that smelters 1-8 are in operation?

16. Unstable Cycles Growing Into a Limit Cycle

Figure 19.11 shows an outward growing spiral that encounters limits after only one or two cycles. These cycles grow rather quickly into a limit cycle, so it may be hard to see the transition from growing oscillations to a limit cycle. To gain a clearer picture of this pattern, repeat the simulation with the initial inventory set just slightly higher than the 3.4 mmt equilibrium value shown in Figure 19.3. How does your new scatter graph compare with Figure 19.11?

17. Is Lumpiness Responsible for Higher Volatility?

Expand the *smelter* dimension to run from 1 to 32. Change the equations for the model in Figure 19.12 to allow the variable cost to run from 35 to 70 cents/pound. Set the variable costs at each smelter to correspond to the shape of the industry cost curve in Figure 19.2, and don't forget to cut the size of the standard smelter in half. Run the new model and compare the results with with Figure 19.14. Does lumpiness contribute to the simulated price volatility?

18. The Cobweb

The final exercise deals with the cobweb theory of cyclical behavior. The cobweb model is a graphical approach explained by Samuelson (1964, p. 396) and Meadows (1980, p. 13) and illustrated for the world aluminum industry in the diagram below.

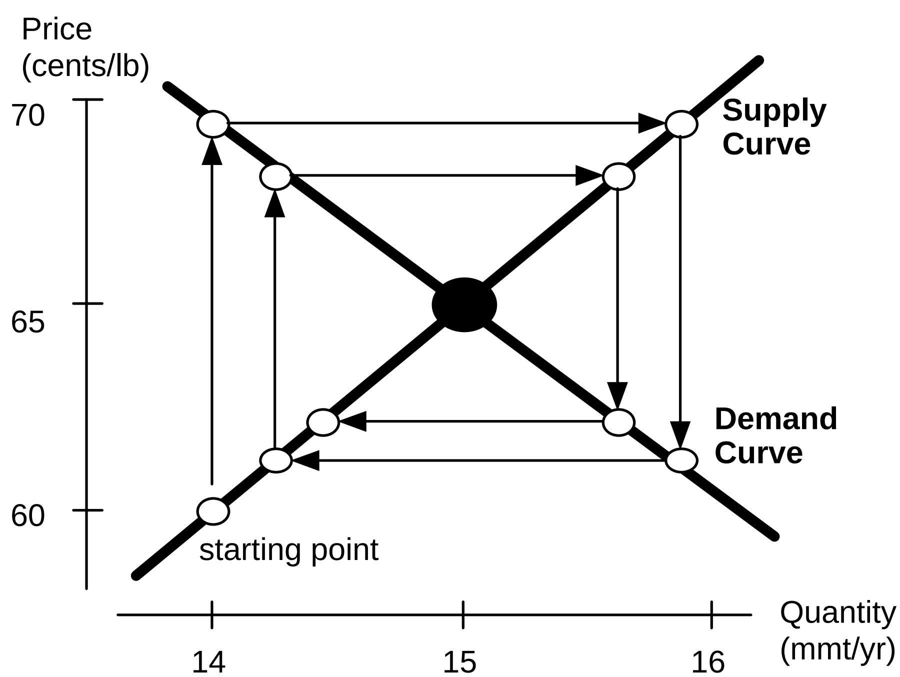

The dark lines show a linear supply curve and a linear demand curve. The intersection point is the market clearing price and quantity. In this case, we expect ingots to sell for 65 cents/pound, and we expect supply and demand to balance at 15 mmt/yr. The slope of the supply curve may be described in terms of the "supply elasticity" which is found by:

relative change in Q:	1/15	= 6.7%
relative change in P:	6/65	= 9.2%
supply elasticity:	6.7%/9.2%	= 0.73

The downward slope of the demand curve may be described in a similar manner. In this case, a price increase of around 5.5 cents/pound is expected to lower the demand by 1 mmt/yr:

relative change in Q:	1/15	= 6.7%
relative change in P:	5.5/65	= 8.5%
demand elasticity:	6.7%/8.5%	= -0.79

The value of -0.79 turns out to be quite close to the long run price elasticity observed in the construction sector.

Now, to begin the graphical analysis, imagine that producers are unable to deliver the market clearing quantity of 15 mmt/yr. For example, a strike may limit their production to 14 mmt/yr, the point on the lower left of the supply curve. Now follow the upward arrow to see the market clearing price of 69 cents/pound. Now follow the arrow to the right to see the quantity of ingots that would be produced in the following time period. If the price were 69 cents/pound, producers would produce around 15.8 mmt/yr. Follow the downward arrow, and you will learn that the market clearing price in the next time period would be around 62 cents/pound.

The graphical analysis shows the arrows leading to the center of the diagram. This implies that the interplay of supply and demand will lead to converging oscillations. But you can tell that the graphical analysis might also have delivered sustained oscillations or even divergent oscillations because slight change in the slopes could deliver different results. For example, we would expect sustained oscillations if the elasticities are identical. The cobweb would yield divergent oscillations if we imagine a flatter supply curve.

With this background, you should be able to answer the following questions about the system dynamics model of aluminum production:

- where does the slope of the supply curve appear in the model?
- where does the slope of the demand curve appear in the model?
- how is the intersection of the two curves found with the model?

The simulations in Figure 19.10 suggests that introducing consumer response to price changes causes greater volatility in the system. Would you have expected this result from the cobweb approach?