

conomy, you'll also see that the
ffect the Federal Reserve Bank.
rowth causes poverty, you'll ask
rowth.

THIS:

B also causes A?

ic world, but a dynamic one.
ad you'll start asking, "What's
up the idea that a system can

ie kind of feedback loop at a
s rarely come singly. They are
ex patterns. A single stock is
ig loops of differing strengths
ay be adjusted by the contents
e stock while it drains another
. The many feedback loops in
e stocks grow, die off, or come
plex systems do much more
pproach goals smoothly—as

— TWO —

A Brief Visit to the Systems Zoo

The . . . goal of all theory is to make the . . . basic elements as simple
and as few as possible without having to surrender the adequate
representation of . . . experience.

—Albert Einstein,¹ physicist

One good way to learn something new is through specific examples rather
than abstractions and generalities, so here are several common, simple
but important examples of systems that are useful to understand in their
own right and that will illustrate many general principles of complex
systems.

This collection has some of the same strengths and weaknesses as a
zoo.² It gives you an idea of the large variety of systems that exist in
the world, but it is far from a complete representation of that variety. It
groups the animals by family—monkeys here, bears there (single-stock
systems here, two-stock systems there)—so you can observe the charac-
teristic behaviors of monkeys, as opposed to bears. But, like a zoo, this
collection is too neat. To make the animals visible and understandable,
it separates them from each other and from their normal concealing
environment. Just as zoo animals more naturally occur mixed together
in ecosystems, so the systems animals described here normally connect
and interact with each other and with others not illustrated here—
all making up the buzzing, hooting, chirping, changing complexity in
which we live.

Ecosystems come later. For the moment, let's look at one system animal
at a time.

One-Stock Systems

A Stock with Two Competing Balancing Loops—a Thermostat

You already have seen the “homing in” behavior of the goal-seeking balancing feedback loop—the coffee cup cooling. What happens if there are two such loops, trying to drag a single stock toward two different goals?

One example of such a system is the thermostat mechanism that regulates the heating of your room (or cooling, if it is connected to an air conditioner instead of a furnace). Like all models, the representation of a thermostat in Figure 15 is a simplification of a real home heating system.




Figure 15. Room temperature regulated by a thermostat and furnace.

Whenever the room temperature falls below the thermostat setting, the thermostat detects a discrepancy and sends a signal that turns on the heat flow from the furnace, warming the room. When the room temperature rises again, the thermostat turns off the heat flow. This straightforward, stock-maintaining, balancing feedback loop is shown on the left side of Figure 15. If there were nothing else in the system, and if you start with a cold room with the thermostat set at 18°C (65°F), it would behave as shown in Figure 16. The furnace comes on, and the room warms up. When the room temperature reaches the thermostat setting, the furnace goes off, and the room stays right at the target temperature.


However, this is not the only loop in the system. Heat also leaks to the outside. The outflow of heat is governed by the second balancing feedback loop, shown on the right side of Figure 15. It is always trying to make the room temperature equal to the outside, just like a coffee cup cooling. If

Figur

Figur

this
show
Th
leak
the s
No
Assu
the
room

f the goal-seeking balance happens if there are two different goals? at mechanism that regulates it is connected to an air s, the representation of a al home heating system.



e thermostat setting, the al that turns on the heat n the room temperature w. This straightforward, own on the left side of n, and if you start with °F), it would behave as e room warms up. When ng, the furnace goes off, .
 1. Heat also leaks to the ond balancing feedback ways trying to make the a coffee cup cooling. If

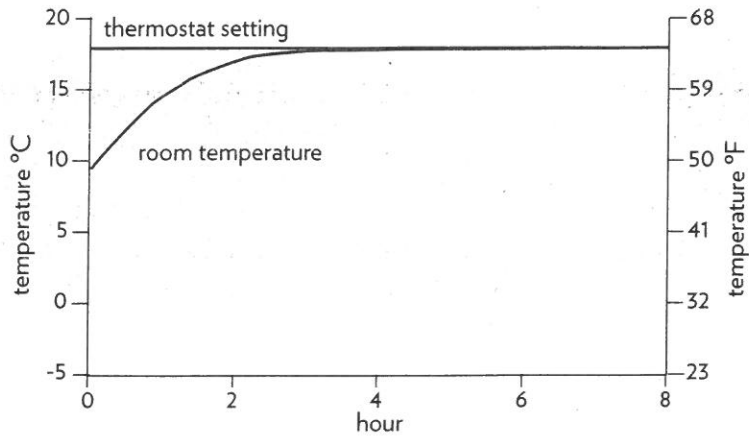


Figure 16. A cold room warms quickly to the thermostat setting.




Figure 17. A warm room cools very slowly to the outside temperature of 10°C.

this were the only loop in the system (if there were no furnace), Figure 17 shows what would happen, starting with a warm room on a cold day.

The assumption is that room insulation is not perfect, and so some heat leaks out of the warm room to the cool outdoors. The better the insulation, the slower the drop in temperature would be.

Now, what happens when these two loops operate at the same time? Assuming that there is sufficient insulation and a properly sized furnace, the heating loop dominates the cooling loop. You end up with a warm room (see Figure 18), even starting with a cold room on a cold day.





Figure 18. The furnace warms a cool room, even as heat continues to leak from the room.

As the room heats up, the heat flowing out of it increases, because there's a larger gap between inside and outside temperatures. But the furnace keeps putting in more heat than the amount that leaks out, so the room warms nearly to the target temperature. At that point, the furnace cycles off and on as it compensates for the heat constantly flowing out of the room.

The thermostat is set at 18°C (65°F) in this simulation, but the room temperature levels off slightly below 18°C (65°F). That's because of the leak to the outside, which is draining away some heat even as the furnace is getting the signal to put it back. This is a characteristic and sometimes surprising behavior of a system with competing balancing loops. It's like trying to keep a bucket full when there's a hole in the bottom. To make things worse, water leaking out of the hole is governed by a feedback loop; the more water in the bucket, the more the water pressure at the hole increases, so the flow out increases! In this case, we are trying to keep the room warmer than the outside and the warmer the room is, the faster it loses heat to the outside. It takes time for the furnace to correct for the increased heat loss—and in that minute still more heat leaks out. In a well-insulated house, the leak will be slower and so the house more comfortable than a poorly insulated one, even a poorly insulated house with a big furnace.

With home heating systems, people have learned to set the thermostat slightly higher than the actual temperature they are aiming at. Exactly how much higher can be a tricky question, because the outflow rate is higher on cold days than on warm days. But for thermostats this control problem



is to leak from the room.

increases, because there's
furnace. But the furnace
leaks out, so the room
temperature, the furnace cycles off
cooling out of the room.
regulation, but the room
temperature is because of the leak
is the furnace is getting
sometimes surprising
It's like trying to keep
the things worse, water
level; the more water in
the room warmer than the
room heat to the outside.
desired heat loss—and in
a poorly insulated one,

to set the thermostat
temperature. Exactly how
outflow rate is higher
this control problem

isn't serious. You can muddle your way to a thermostat setting you can live with.

For other systems with this same structure of competing balancing loops, the fact that the stock goes on changing while you're trying to control it can create real problems. For example, suppose you're trying to maintain a store inventory at a certain level. You can't instantly order new stock to make up an immediately apparent shortfall. If you don't account for the goods that will be sold while you are waiting for the order to come in, your inventory will never be quite high enough. You can be fooled in the same way if you're trying to maintain a cash balance at a certain level, or the level of water in a reservoir, or the concentration of a chemical in a continuously flowing reaction system.

There's an important general principle here, and also one specific to the thermostat structure. First the general one: The information delivered by a feedback loop can only affect future behavior; it can't deliver the information, and so can't have an impact fast enough to correct behavior that drove the current feedback. A person in the system who makes a decision based on the feedback can't change the behavior of the system that drove the current feedback; the decisions he or she makes will affect only future behavior. *Causality*

The information delivered by a feedback loop—even nonphysical feedback—can only affect future behavior; it can't deliver a signal fast enough to correct behavior that drove the current feedback. Even nonphysical information takes time to feedback into the system.

Why is that important? Because it means there will always be delays in responding. It says that a flow can't react instantly to a flow. It can react only to a change in a stock, and only after a slight delay to register the incoming information. In the bathtub, it takes a split second of time to assess the depth of the water and decide to adjust the flows. Many economic models make a mistake in this matter by assuming that consumption or production can respond immediately, say, to a change in price. That's one of the reasons why real economies tend not to behave exactly like many economic models.

The specific principle you can deduce from this simple system is that you must remember in thermostat-like systems to take into account whatever draining or filling processes are going on. If you don't, you won't achieve the target level of your stock. If you want your room temperature to be at 18°C (65°F), you have to set the thermostat a little above the desired

temperature. If you want to pay off your credit card (or the national debt), you have to raise your repayment rate high enough to cover the charges you incur while you're paying (including interest). If you're gearing up

A stock-maintaining balancing feedback loop must have its goal set appropriately to compensate for draining or inflowing processes that affect that stock. Otherwise, the feedback process will fall short of or exceed the target for the stock.

your work force to a higher level, you have to hire fast enough to correct for those who quit while you are hiring. In other words, your mental model of the system needs to include all the important flows, or you will be surprised by the system's behavior.

Before we leave the thermostat, we should see how it behaves with a varying outside temperature. Figure 19 shows a twenty-four-hour period of normal operation of a well-functioning thermostat system, with the outside temperature dipping well below freezing. The inflow of heat from the furnace nicely tracks the outflow of heat to the outside. The temperature in the room varies hardly at all once the room has warmed up.

Every balancing feedback loop has its breakdown point, where other loops pull the stock away from its goal more strongly than it can pull back. That can happen in this simulated thermostat system, if I weaken the power of the heating loop (a smaller furnace that cannot put out as much heat), or if I strengthen the power of the cooling loop (colder outside tempera-




Figure 19. The furnace warms a cool room, even as heat leaks from the room and outside temperatures drop below freezing.

ture,
the s:
from
with
down
room

temperature °C

Figure


See
relate
tures a
outside
enough
and th
But
furnac
ating l
outside
operat
again.

Just
more l
the inf
study t

l (or the national debt),
gh to cover the charges
) . If you're gearing up
gher level, you have to
ect for those who quit
In other words, your
em needs to include all
ou will be surprised by

rmostat, we should see
ying outside tempera-
enty-four-hour period
well-functioning ther-
ig well below freezing.
: outflow of heat to the
v at all once the room

n point, where other
v than it can pull back.
, if I weaken the power
ut out as much heat),
lder outside tempera-



om the room and outside

ture, less insulation, or larger leaks). Figure 20 shows what happens with the same outside temperatures as in Figure 19, but with faster heat loss from the room. At very cold temperatures, the furnace just can't keep up with the heat drain. The loop that is trying to bring the room temperature down to the outside temperature dominates the system for a while. The room gets pretty uncomfortable!




Figure 20. On a cold day, the furnace can't keep the room warm in this leaky house!

See if you can follow, as time unfolds, how the variables in Figure 20 relate to one another. At first, both the room and the outside temperatures are cool. The inflow of heat from the furnace exceeds the leak to the outside, and the room warms up. For an hour or two, the outside is mild enough that the furnace replaces most of the heat that's lost to the outside, and the room temperature stays near the desired temperature.

But as the outside temperature falls and the heat leak increases, the furnace cannot replace the heat fast enough. Because the furnace is generating less heat than is leaking out, the room temperature falls. Finally, the outside temperature rises again, the heat leak slows, and the furnace, still operating at full tilt, finally can pull ahead and start to warm the room again.

Just as in the rules for the bathtub, whenever the furnace is putting in more heat than is leaking out, the room temperature rises. Whenever the inflow rate falls behind the outflow rate, the temperature falls. If you study the system changes on this graph for a while and relate them to the

feedback-loop diagram of this system, you'll get a good sense of how the structural interconnections of this system—its two feedback loops and how they shift in strength relative to each other—lead to the unfolding of the system's behavior over time.

A Stock with One Reinforcing Loop and One Balancing Loop—Population and Industrial Economy

What happens when a reinforcing and a balancing loop are both pulling on the same stock? This is one of the most common and important system structures. Among other things, it describes every living population and every economy.

A population has a reinforcing loop causing it to grow through its birth rate, and a balancing loop causing it to die off through its death rate.




Figure 21. Population governed by a reinforcing loop of births and a balancing loop of deaths.

As long as fertility and mortality are constant (which in real systems they rarely are), this system has a simple behavior. It grows exponentially or dies off, depending on whether its reinforcing feedback loop determining births is stronger than its balancing feedback loop determining deaths.

For example, the 2007 world population of 6.6 billion people had a fertility rate of roughly 21 births a year for every 1,000 people in the population. Its mortality rate was 9 deaths a year out of every 1,000 people. Fertility was higher than mortality, so the reinforcing loop dominated the system. If those fertility and mortality rates continue unchanged, a child born now will see the world population more than double by the time he or she reaches the age of 60, as shown in Figure 22.

If, because of a terrible disease, the mortality rate were higher, say at 30

population (billions)

Figure 2:
nine dea

deaths would be bor

population (billions)


Figure 2:
is much h

Thin
time. V

a good sense of how the two feedback loops and lead to the unfolding of

Population and
g loop are both pulling
n and important system
y living population and

o grow through its birth
ugh its death rate.



nd a balancing loop of deaths.

(which in real systems
It grows exponentially
g feedback loop deter-
ack loop determining

lion people had a fertil-
ople in the population.
1,000 people. Fertility
dominated the system.
changed, a child born
le by the time he or she

e were higher, say at 30




Figure 22. Population growth if fertility and mortality maintain their 2007 levels of 21 births and nine deaths per 1,000 people.

deaths per 1,000, while the fertility rate remained at 21, then the death loop would dominate the system. More people would die each year than would be born, and the population would gradually decrease (Figure 23).




Figure 23. Population decline if fertility remains at 2007 level (21 births per 1,000) but mortality is much higher, 30 deaths per 1,000.

Things get more interesting when fertility and mortality change over time. When the United Nations makes long-range population projections,

it generally assumes that, as countries become more developed, average fertility will decline (approaching replacement where on average each woman has 1.85 children). Until recently, assumptions about mortality were that it would also decline, but more slowly (because it is already low in most parts of the world). However, because of the epidemic of HIV/AIDS, the UN now assumes the trend of increasing life expectancy over the next fifty years will slow in regions affected by HIV/AIDS.

Changing flows (fertility and mortality) create a change in the behavior over time of the stock (population)—the line bends. If, for example, world fertility falls steadily to equal mortality by the year 2035 and they both stay constant thereafter, the population will level off, births exactly balancing deaths in dynamic equilibrium, as in Figure 24.




Figure 24. Population stabilizes when fertility equals mortality.


This behavior is an example of **shifting dominance** of feedback loops. Dominance is an important concept in systems thinking. When one loop dominates another, it has a stronger impact on behavior. Because systems often have several competing feedback loops operating simultaneously, those loops that dominate the system will determine the behavior.

At first, when fertility is higher than mortality, the reinforcing growth loop dominates the system and the resulting behavior is exponential growth. But that loop is gradually weakened as fertility falls. Finally, it exactly equals the strength of the balancing loop of mortality. At that point neither loop dominates, and we have dynamic equilibrium.

You
tempe
overw
heat in
heating
They
could l
to the
ity. The
system
loop. A
expon
die off
off if t
sequen
two loc
I cho
about
confron
econon
change
ing), th
good a

The first
and the
opinion
ally hap
happens

ore developed, average here on average each tions about mortality ecause it is already low the epidemic of HIV/ life expectancy over the AIDS. change in the behavior . If, for example, world 035 and they both stay irths exactly balancing



nce of feedback loops. nking. When one loop avior. Because systems rating simultaneously, e the behavior. he reinforcing growth avior is exponential rtility falls. Finally, it mortality. At that point ibrrium.

You saw shifting dominance in the thermostat system, when the outside temperature fell and the heat leaking out of the poorly insulated house overwhelmed the ability of the furnace to put heat into the room. Dominance shifted from the heating loop to the cooling loop.

Complex behaviors of systems often arise as the relative strengths of feedback loops shift, causing first one loop and then another to dominate behavior.

There are only a few ways a population system could behave, and these depend on what happens to the “driving” variables, fertility and mortality. These are the only ones possible for a simple system of one reinforcing and one balancing loop. A stock governed by linked reinforcing and balancing loops will grow exponentially if the reinforcing loop dominates the balancing one. It will die off if the balancing loop dominates the reinforcing one. It will level off if the two loops are of equal strength (see Figure 25). Or it will do a sequence of these things, one after another, if the relative strengths of the two loops change over time (see Figure 26).

I chose some provocative population scenarios here to illustrate a point about models and the scenarios they can generate. Whenever you are confronted with a scenario (and you are, every time you hear about an economic prediction, a corporate budget, a weather forecast, future climate change, a stockbroker saying what is going to happen to a particular holding), there are questions you need to ask that will help you decide how good a representation of reality is the underlying model.

- Are the driving factors likely to unfold this way? (What are birth rate and death rate likely to do?)
- If they did, would the system react this way? (Do birth and death rates really cause the population stock to behave as we think it will?)
- What is driving the driving factors? (What affects birth rate? What affects death rate?)

The first question can't be answered factually. It's a guess about the future, and the future is inherently uncertain. Although you may have a strong opinion about it, there's no way to prove you're right until the future actually happens. A systems analysis can test a number of scenarios to see what happens if the driving factors do different things. That's usually one purpose

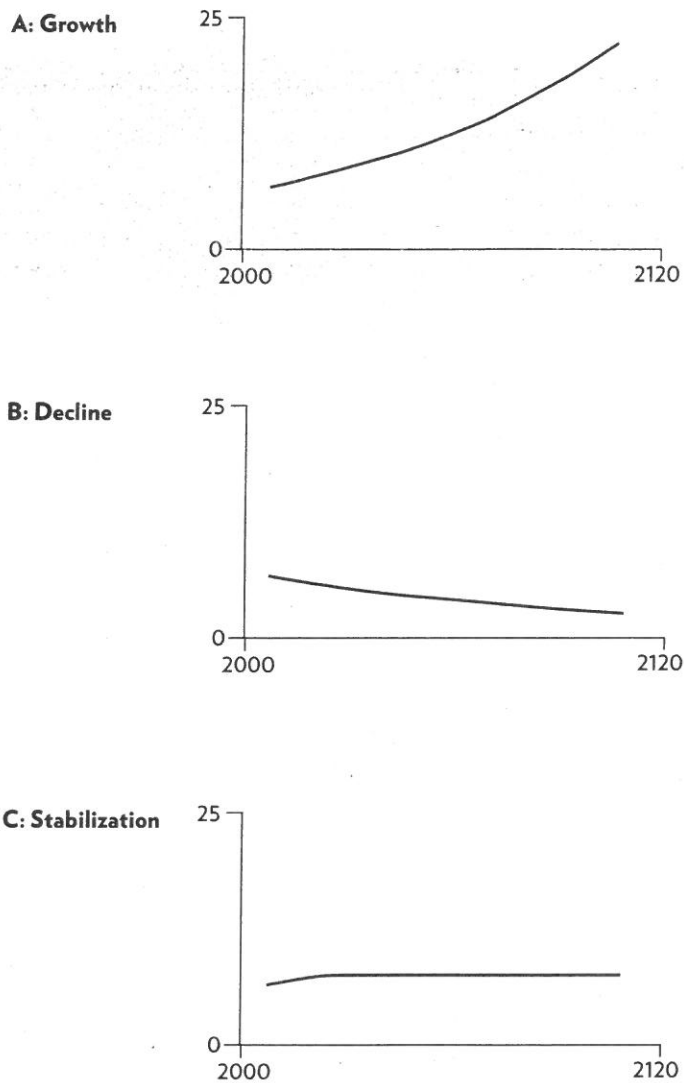


Figure 25. Three possible behaviors of a population: growth, decline, and stabilization.

of a systems analysis. But you have to be the judge of which scenario, if any, should be taken seriously as a future that might really be possible.

Dynamic systems studies usually are not designed to *predict* what will happen. Rather, they're designed to explore *what would happen*, if a number of driving factors unfold in a range of different ways.

Figure 20

The
more sc
the inh
driving

In the
answer
like this
The pop
simple...
age gro
respond
ing und
when a
behavior
Finally




Figure 26. Shifting dominance of fertility and mortality loops.

The second question—whether the system really will react this way—is more scientific. It’s a question about how good the model is. Does it capture the inherent dynamics of the system? Regardless of whether you think the driving factors *will* do that, *would the system behave like that* if they did?

In the population scenarios above, however likely you think they are, the answer to this second question is roughly yes, the population would behave like this, if the fertility and mortality did that.

The population model I have used here is very simple. A more detailed model would distinguish age groups, for example. But basically this model responds the way a real population would, growing under the conditions when a real population would grow, declining when a real population would decline. The numbers are off, but the basic behavior pattern is realistic.

System dynamics models explore possible futures and ask “what if” questions.

Finally, there is the third question. What is driving the driving factors?

QUESTIONS FOR TESTING THE VALUE OF A MODEL

1. Are the driving factors likely to unfold this way?
2. If they did, would the system react this way?
3. What is driving the driving factors?

decline, and stabilization.

e of which scenario, if any, really be possible. gned to *predict* what will *would happen*, if a number ways.

Model utility depends not on whether its driving scenarios are realistic (since no one can know that for sure), but on whether it responds with a realistic pattern of behavior.

What is adjusting the inflows and outflows? This is a question about system boundaries. It requires a hard look at those driving factors to see if they are actually independent, or if they are also embedded in the system.

Is there anything about the size of the population, for instance, that might feed back to influence fertility or mortality? Do other factors—economics, the environment, social trends—influence fertility and mortality? Does the size of the population affect those economic and environmental and social factors?

Of course, the answer to all of these questions is yes. Fertility and mortality are governed by feedback loops too. At least some of those feedback loops are themselves affected by the size of the population. This population “animal” is only one piece of a much larger system.³

One important piece of the larger system that affects population is the economy. At the heart of the economy is another reinforcing-loop-plus-balancing-loop system—the same kind of structure, with the same kinds of behavior, as the population (see Figure 27).

The greater the stock of physical capital (machines and factories) in the economy and the efficiency of production (output per unit of capital), the more output (goods and services) can be produced each year.

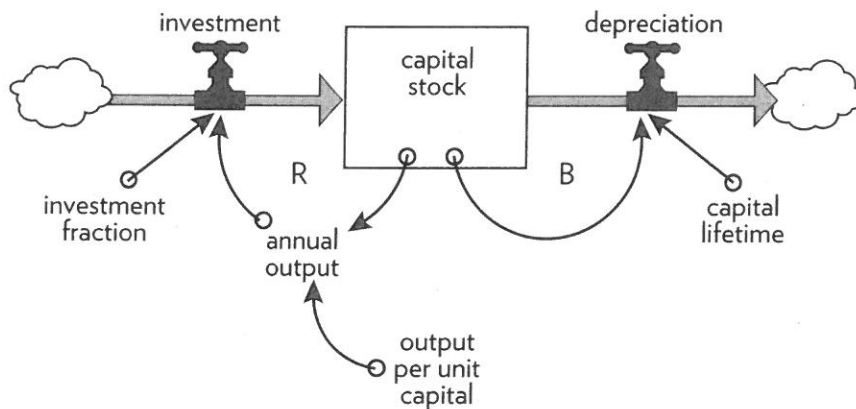


Figure 27. Like a living population, economic capital has a reinforcing loop (investment of output) governing growth and a balancing loop (depreciation) governing decline.

The capital invests out. Phys out. Th loop in age cap capital. If thi have th world p growin dies off than its

If a c efficien capital s of the c capital l than it i econom econom grows ex This i can mak ing its ir Just as tion, so the lifeti


inflows and outflows? At system boundaries. It those driving factors to independent, or if they are stem.

out the size of the population might feed back to mortality? Do other ends—influence fertility affect those economic and

yes. Fertility and mortality—some of those feedback loops. This population n^3

affects population is the reinforcing-loop-plus-re, with the same kinds

ness and factories) in the t per unit of capital), the d each year.



reinforcing loop (investment of governing decline.

The more output that is produced, the more can be invested to make new capital. This is a reinforcing loop, like the birth loop for a population. The investment fraction is equivalent to the fertility. The greater the fraction of its output a society invests, the faster its capital stock will grow.

Physical capital is drained by depreciation—obsolescence and wearing-out. The balancing loop controlling depreciation is equivalent to the death loop in a population. The “mortality” of capital is determined by the average capital lifetime. The longer the lifetime, the smaller the fraction of capital that must be retired and replaced each year.

If this system has the same structure as the population system, it must have the same repertoire of behaviors. Over recent history world capital, like world population, has been dominated by its reinforcing loop and has been growing exponentially. Whether in the future it grows or stays constant or dies off depends on whether its reinforcing growth loop remains stronger than its balancing depreciation loop. That depends on

- the investment fraction—how much output the society invests rather than consumes,
- the efficiency of capital—how much capital it takes to produce a given amount of output, and
- the average capital lifetime.

If a constant fraction of output is reinvested in the capital stock and the efficiency of that capital (its ability to produce output) is also constant, the capital stock may decline, stay constant, or grow, depending on the lifetime of the capital. The lines in Figure 28 show systems with different average capital lifetimes. With a relatively short lifetime, the capital wears out faster than it is replaced. Reinvestment does not keep up with depreciation and the economy slowly declines. When depreciation just balances investment, the economy is in dynamic equilibrium. With a long lifetime, the capital stock grows exponentially. The longer the lifetime of capital, the faster it grows.

This is another example of a principle we’ve already encountered: You can make a stock grow by decreasing its outflow rate as well as by increasing its inflow rate!

Just as many factors influence the fertility and mortality of a population, so many factors influence the output ratio, investment fraction, and the lifetime of capital—interest rates, technology, tax policy, consumption




Figure 28. Growth in capital stock with changes in the lifetime of the capital. In a system with output per unit capital ratio of 1:3 and an investment fraction of 20 percent, capital with a lifetime of 15 years just keeps up with depreciation. A shorter lifetime leads to a declining stock of capital.

habits, and prices, to name just a few. Population itself influences investment, both by contributing labor to output, and by increasing demands on consumption, thereby decreasing the investment fraction. Economic output also feeds back to influence population in many ways. A richer economy usually has better health care and a lower death rate. A richer economy also usually has a lower birth rate.

In fact, just about any long-term model of a real economy should link together the two structures of population and capital to show how they affect each other. The central question of economic development is how to keep the reinforcing loop of capital accumulation from growing more slowly than the reinforcing loop of population growth—so that people are getting richer instead of poorer.⁴

It may seem strange to you that I call the capital system the same kind of “zoo animal” as the population system. A production system with factories

Systems with similar feedback structures produce similar dynamic behaviors.

and shipments and economic flows doesn't look much like a population system with babies being born and people aging and having more babies and dying. But from a systems point of view these systems, so dissimilar in many ways, have one

important thing in common: their feedback-loop structures. Both have a stock governed by a reinforcing growth loop and a balancing death loop.


Both
older
On
tion
simil
outw
A p
repro
both
like a
econo
back l

A Syste
Pictur
delive
of cars
Now
tory h
Figure



oi
to f

Figure 2:
through s



time of the capital. In a system with a return on investment of 20 percent, capital with a shorter lifetime leads to a declining

on itself influences investment by increasing demands on investment fraction. Economic growth in many ways. A richer economy has a lower death rate. A richer

real economy should link investment capital to show how they influence economic development is how investment from growing more economic growth—so that people are

in a real system the same kind of investment system with factories and economic flows doesn't look like a system with babies being born and having more babies. From this point of view these systems in many ways, have one of the same structures. Both have a balancing death loop.

Both also have an aging process. Steel mills and lathes and turbines get older and die just as people do.

One of the central insights of systems theory, as central as the observation that systems largely cause their own behavior, is that systems with similar feedback structures produce similar dynamic behaviors, even if the outward appearance of these systems is completely dissimilar.

A population is nothing like an industrial economy, except that both can reproduce themselves out of themselves and thus grow exponentially. And both age and die. A coffee cup cooling is like a warmed room cooling, and like a radioactive substance decaying, and like a population or industrial economy aging and dying. Each declines as the result of a balancing feedback loop.

A System with Delays—Business Inventory

Picture a stock of inventory in a store—a car dealership—with an inflow of deliveries from factories and an outflow of new car sales. By itself, this stock of cars on the dealership lot would behave like the water in a bathtub.

Now picture a regulatory feedback system designed to keep the inventory high enough so that it can always cover ten days' worth of sales (see Figure 29). The car dealer needs to keep some inventory because deliveries




Figure 29. Inventory at a car dealership is kept steady by two competing balancing loops, one through sales and one through deliveries.

and purchases don't match perfectly every day. Customers make purchases that are unpredictable on a day-to-day basis. The car dealer also needs to provide herself with some extra inventory (a buffer) in case deliveries from suppliers are delayed occasionally.

The dealer monitors sales (perceived sales), and if, for example, they seem to be rising, she adjusts orders to the factory to bring inventory up to her new desired inventory that provides ten days' coverage at the higher sales rate. So, higher sales mean higher perceived sales, which means a higher discrepancy between inventory and desired inventory, which means higher orders, which will bring in more deliveries, which will raise inventory so it can comfortably supply the higher rate of sales.

This system is a version of the thermostat system—one balancing loop of sales draining the inventory stock and a competing balancing loop maintaining the inventory by resupplying what is lost in sales. Figure 30 shows the not very surprising result of an increase in consumer demand of 10 percent.

In Figure 31, I am putting something else into this simple model—three delays that are typical of what we experience in the real world.

First, there is a perception delay, intentional in this case. The car dealer doesn't react to just any blip in sales. Before she makes ordering decisions, she averages sales over the past five days to sort out real trends from temporary dips and spikes.

Second, there is a response delay. Even when it's clear that orders need




Figure 30. Inventory on the car dealership's lot with a permanent 10-percent increase in consumer demand starting on day 25.




Figure 31
perception

to be ad
order. B
Another
days to l
five day:
deliver i
Altho
the simp
system.

cars on the lot

Figure 32.
system.

Customers make purchases (The car dealer also needs to offer) in case deliveries from

and if, for example, they seem to bring inventory up to her coverage at the higher sales rates, which means a higher inventory, which means higher price will raise inventory so it

them—one balancing loop of setting balancing loop maintains in sales. Figure 30 shows a consumer demand of 10

to this simple model—three the real world.

in this case. The car dealer makes ordering decisions, but real trends from tempo-

it's clear that orders need

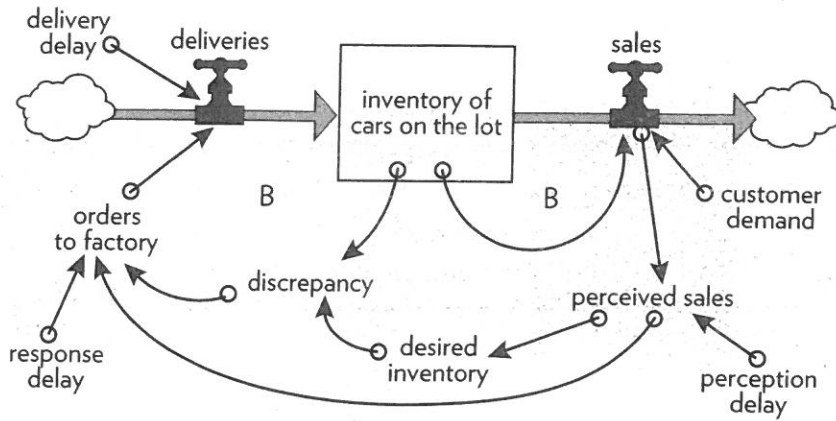


Figure 31. Inventory at a car dealership with three common delays now included in the picture—a perception delay, a response delay, and a delivery delay.

to be adjusted, she doesn't try to make up the whole adjustment in a single order. Rather, she makes up one-third of any shortfall with each order. Another way of saying that is, she makes partial adjustments over three days to be extra sure the trend is real. Third, there is a delivery delay. It takes five days for the supplier at the factory to receive an order, process it, and deliver it to the dealership.

Although this system still consists of just two balancing loops, like the simplified thermostat system, it doesn't behave like the thermostat system. Look at what happens, for example, as shown in Figure 32, when





Figure 32. Response of inventory to a 10-percent increase in sales when there are delays in the system.



permanent 10-percent increase in

the business experiences the same permanent 10-percent jump in sales from an increase in customer demand.

Oscillations! A single step up in sales causes inventory to drop. The car dealer watches long enough to be sure the higher sales rate is going to last. Then she begins to order more cars to both cover the new rate of sales and bring the inventory up. But it takes time for the orders to come in. During that time inventory drops further, so orders have to go up a little more, to bring inventory back up to ten days' coverage.

Eventually, the larger volume of orders starts arriving, and inventory recovers—and more than recovers, because during the time of uncertainty about the actual trend, the owner has ordered too much. She now sees her mistake, and cuts back, but there are still high past orders coming in, so she orders even less. In fact, almost inevitably, since she still can't be sure of what is going to happen next, she orders too little. Inventory gets too low again. And so forth, through a series of oscillations around the new desired inventory level. As Figure 33 illustrates, what a difference a few delays make!

We'll see in a moment that there are ways to damp these oscillations in inventory, but first it's important to understand why they occur. It isn't because the car dealer is stupid. It's because she is struggling to operate in a system in which she doesn't have, and can't have, timely information and in which physical delays prevent her actions from having an immediate effect on inventory. She doesn't know what her customers will do

A delay in a balancing feedback loop makes a system likely to oscillate.

next. When they do something, she's not sure they'll keep doing it. When she issues an order, she doesn't see an immediate response. This situation of information insufficiency and physical delays is very common. Oscillations like these are frequently encountered in inventories and in many other systems. Try taking a shower sometime where there's a very long pipe between the hot- and cold-water mixer and the showerhead, and you'll experience directly the joys of hot and cold oscillations because of a long response delay.

How much of a delay causes what kind of oscillation under what circumstances is not a simple matter. I can use this inventory system to show you why.

"These oscillations are intolerable," says the car dealer (who is herself a learning system, determined now to change the behavior of the inventory

Figure
but shai
the cha
from th

system
the de
averag
adjust

Figu
shorte
Not
If any
worse.

cars per day

10-percent jump in sales

inventory to drop. The car sales rate is going to last. For the new rate of sales and orders to come in. During the time to go up a little more, to

is arriving, and inventory is dropping during the time of uncertainty is too much. She now sees her past orders coming in, so since she still can't be sure about the future, a little. Inventory gets too low, and oscillations around the new level, what a difference a few

to damp these oscillations in inventory and why they occur. It isn't just the dealer who is struggling to operate. The dealer, to have, timely information is needed from having an immediate response from her customers will do something, she's not sure when she issues an order, she

response. This situation is a result of the delay in the ordering process and physical delays in the system. It is like these are frequently observed in many systems. Try taking a shower with the hot- and cold-water taps. You can feel directly the joys of hot water and the pains of cold water. The delay in the water temperature is a direct result of the delay in the system.

under what circumstances a control system to show you how the system behaves. The car dealer (who is herself a control system) is a good example of the behavior of the inventory

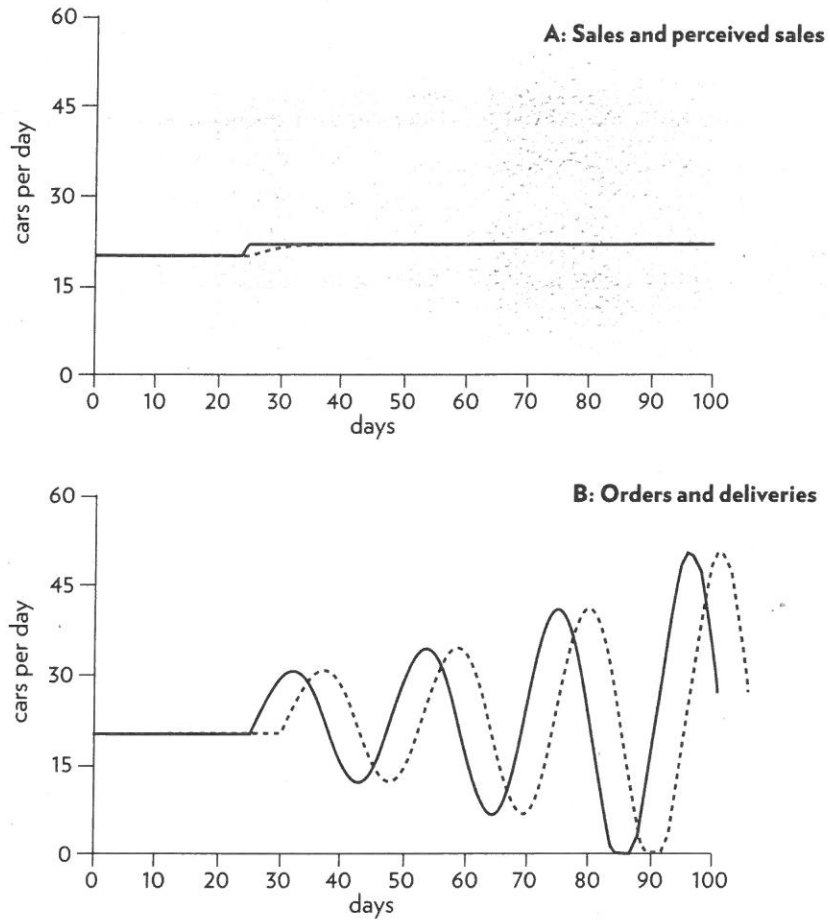


Figure 33. The response of orders and deliveries to an increase in demand. A shows the small but sharp step up in sales on day 25 and the car dealer's "perceived" sales, in which she averages the change over 3 days. B shows the resulting ordering pattern, tracked by the actual deliveries from the factory.

system). "I'm going to shorten the delays. There's not much I can do about the delivery delay from the factory, so I'm going to react faster myself. I'll average sales trends over only two days instead of five before I make order adjustments."

Figure 34 illustrates what happens when the dealer's perception delay is shortened from five days to two.

Not much happens when the car dealer shortens her perception delay. If anything the oscillations in the inventory of cars on the lot are a bit worse. And if, instead of shortening her perception time, the car dealer

tries shortening her reaction time—making up perceived shortfalls in two days instead of three—things get very much worse, as shown in Figure 35.

Something has to change and, since this system has a learning person within it, something will change. “High leverage, wrong direction,” the system-thinking car dealer says to herself as she watches this failure of a policy intended to stabilize the oscillations. This perverse kind of result can




Figure 34. The response of inventory to the same increase in demand with a shortened perception delay.




Figure 35. The response of inventory to the same increase in demand with a shortened reaction time. Acting faster makes the oscillations worse!

be see
to a p
then t
is just
behavi


Part
slowly,
overre
delay fi
to six, a
As Fi
and the
efficien
The 1
the one
the car
factory.
that par
manage

Chang
much e
can see v
We're al


cars

Figure 36.
time.

perceived shortfalls in two
se, as shown in Figure 35.
m has a learning person
ge, wrong direction," the
watches this failure of a
erverse kind of result can



emand with a shortened percep-



emand with a shortened reaction

be seen all the time—someone trying to fix a system is attracted intuitively to a policy lever that in fact does have a strong effect on the system. And then the well-intentioned fixer pulls the lever in the wrong direction! This is just one example of how we can be surprised by the counterintuitive behavior of systems when we start trying to change them.

Part of the problem here is that the car dealer has been reacting not too slowly, but too quickly. Given the configuration of this system, she has been overreacting. Things would go better if, instead of decreasing her response delay from three days to two, she would increase the delay from three days to six, as illustrated in Figure 36.

As Figure 36 shows, the oscillations are greatly damped with this change, and the system finds its new equilibrium fairly efficiently.

The most important delay in this system is the one that is not under the direct control of the car dealer. It's the delay in delivery from the factory. But even without the ability to change that part of her system, the dealer can learn to manage inventory quite well.

Changing the delays in a system can make it much easier or much harder to manage. You can see why system thinkers are somewhat fanatic on the subject of delays. We're always on the alert to see where delays occur in systems, how long

Delays are pervasive in systems, and they are strong determinants of behavior. **Changing the length of a delay may (or may not, depending on the type of delay and the relative lengths of other delays) make a large change in the behavior of a system.**




Figure 36. The response of inventory to the same increase in demand with a slowed reaction time.

they are, whether they are delays in information streams or in physical processes. We can't begin to understand the dynamic behavior of systems unless we know where and how long the delays are. And we are aware that some delays can be powerful policy levers. Lengthening or shortening them can produce major changes in the behavior of systems.

In the big picture, one store's inventory problem may seem trivial and fixable. But imagine that the inventory is that of all the unsold automobiles in America. Orders for more or fewer cars affect production not only at assembly plants and parts factories, but also at steel mills, rubber and glass plants, textile producers, and energy producers. Everywhere in this system are perception delays, production delays, delivery delays, and construction delays. Now consider the link between car production and jobs—increased production increases the number of jobs allowing more people to buy cars. That's a reinforcing loop, which also works in the opposite direction—less production, fewer jobs, fewer car sales, less production. Put in another reinforcing loop, as speculators buy and sell shares in the auto and auto-supply companies based on their recent performance, so that an upsurge in production produces an upsurge in stock price, and vice versa.

That very large system, with interconnected industries responding to each other through delays, entraining each other in their oscillations, and being amplified by multipliers and speculators, is the primary cause of business cycles. Those cycles don't come from presidents, although presidents can do much to ease or intensify the optimism of the upturns and the pain of the downturns. Economies are extremely complex systems; they are full of balancing feedback loops with delays, and they are inherently oscillatory.⁵

Two-Stock Systems

A Renewable Stock Constrained by a Nonrenewable Stock—an Oil Economy

The systems I've displayed so far have been free of constraints imposed by their surroundings. The capital stock of the industrial economy model didn't require raw materials to produce output. The population didn't need food. The thermostat-furnace system never ran out of oil. These simple models of the systems have been able to operate according to their unconstrained internal dynamics, so we could see what those dynamics are.

But any real physical entity is always surrounded by and exchanging things

with i
mater
needs
food :
jobs a
entity
waste:

The
const
loop t
the gr
the in

Gro
system
more
iar bel
wheth
epider

are dr
We kn
ing th
system
will se
tion ir
of fuel
infect.
capita
manag

Like
a poll
newab
no cap
enviro
said h
same
system

The