Economics from the Top Down

new ideas in economics and the social sciences

Insights from the Lotka-Volterra Model

Blair Fix

September 15, 2025

All models are wrong;
the practical question is how wrong do they have to be to not be useful.

— George Box

In science, there's an inherent trade off between comprehensibility and realism. Realistic models tend to be intricate ... even convoluted. But to be comprehensible, a model must be simple.

For a good example of this trade off, look to high-school physics. In the real world, we know that projectiles are affected by aerodynamics. (That's why frisbees fly differently than baseballs.) But since aerodynamics are complicated, high school teachers ignore them. Instead, they teach students that earthbound projectiles behave as they would on the moon — blissfully unaffected by air drag. This simplification is a lie, of course. But it's useful for teaching students about the essence of Newton's equations.

Science is filled with this sort of simplification. We learn about the world by developing toy models — models which simplify reality, yet retain (we hope) an element of truth.

In economics, there's no shortage of toy models. But most of these playthings belong in the landfill; they're models that assume away the most pertinent features of the real world. (For example, neoclassical economic models capitalism by assuming 'perfect competition', whereas the real world is marked by pernicious oligarchy.)

In short, if we want simple models that capture key elements of human behavior, it's best to leave mainstream economics behind. Instead, a good place to start is with population biology — specifically the Lotka-Volterra model of predator-prey dynamics. Like projectile motion that neglects aerodynamics, the Lotka-Volterra equations are a toy model of how predator and prey populations respond to each other. In a sense, it's the simplest 'systems model' that still provides useful insights about the real world.

In what follows, we'll take a tour of the Lotka-Volterra model, and see how it gives insights into human behavior.

The Lotka-Volterra model

Developed in the early 20th century by the mathematicians Alfred J. Lotka and Vito Volterra, the Lotka-Volterra equations are an early example of what we would today call a 'systems model' — a model that simulates feedback between two or more entities.

In the Lotka-Volterra model, we imagine feedback between a population of predators and a population of prey. Looking ahead, an important feature of the Lotka-Volterra equations is that they can't be solved with algebra. Instead, they must be solved numerically by throwing numbers in and seeing what comes out.

Why is this algebraic intractability important? Well, because much of mainstream economics operates like a subdiscipline of pure mathematics, where the goal is to postulate models with neat (but worthless) analytic solutions. And since systems models defy this sort of rigid thinking, they've been neglected by economists.

Back to the Lotka-Volterra model. The model begins by imagining two populations, one of prey and one of predators. Now the presumption is that predators eat prey. However this predation isn't captured literally by the model. (There's no equation that tells us when or how a wolf eats a sheep.) Instead, the model simulates predation in terms of population dynamics. For example, if a wolf population expands, it will cut into the growth rate of the neighboring sheep population (on which it preys).

The Lotka-Volterra model assumes that if left in isolation, our predator and prey populations will have opposite dynamics. If left alone, our prey population will *grow* exponentially.[1] (For example, if we put a group of sheep into an empty field, they will reproduce and their population will expand.) In contrast, if our predator population is left alone it will *decline* exponentially. (For example, if we deprive a wolf pack of food, its population will gradually starve to death.)

So those are the starting assumptions, which both yield straightforward predictions if they're left to play out. Fortunately, the Lotka-Volterra model doesn't stop there. Instead, it imagines what happens when we *mix* predators with prey. It's here that we encounter the magic of feedback. If wolves kill sheep, then a larger wolf population will reduce the growth of the sheep population. But if the sheep population declines, that causes the wolf population to die off from starvation.

Now an economist might look at this model and imagine that it arrives at some sort of equilibrium with an 'optimal' number of sheep and wolves. But that's *not* what happens. When we run the Lotka-Volterra model, we find that it's defined by non-equilibrium dynamics of boom and bust.

Figure 1 illustrates the dance between predators and prey. Here, the blue curve shows the population of sheep, while the red curve shows the population of wolves. Initially, the population of wolves is small enough that the sheep population expands happily. But this sheep boom then causes the wolf population to grow. As the wolf population balloons, predation causes a collapse of the sheep population, leading to starvation amongst the wolves. Finally, once enough wolves have died, the sheep population starts to expand, and the cycle begins again.

Now, the cyclical dance of the Lotka-Volterra model is well known to population biologists. In fact, it illustrates a fundamental feature of natural systems: they're marked not by static equilibrium, but by what the physicist Ilya Prigogine called 'order through fluctuation'. In short, if natural systems are stable, it's *because* they fluctuate. And if they *don't* fluctuate (if they veer in a single direction), that's a sign that something abnormal is happening.

[1] In the real world, the sheep population will eventually plateau as it reaches the field's carrying capacity. But the Lotka-Volterra model assumes that predation keeps the sheep population well below this upper limit.

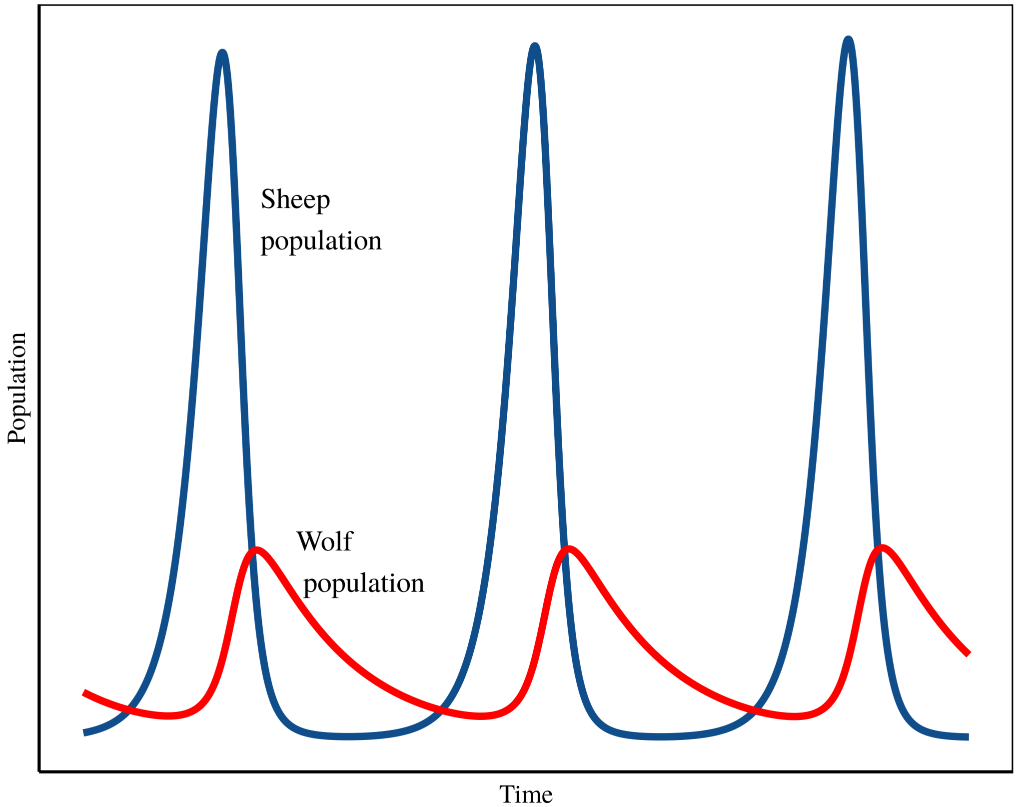

Figure 1: Boom-bust dynamics — the characteristic behavior of the Lotka-Volterra model

A key feature of the Lotka-Volterra model is that it gives rise to cycles of booms and busts, here visualized by a feedback relation between sheep and wolf populations. A rising sheep population causes the wolf population to boom. Over-predation then causes the sheep population to collapse, leading to an eventual decline in the number of wolves. Once enough wolves have died off, the cycle begins again.

Harvesting non-renewable resources

Speaking of 'abnormal', let's think about the nature of industrial society. It's built on a one-time bonanza of fossil-fuel extraction, which means that there can be no long-term cycles. When it comes to fossil fuels, order through fluctuation gives way to a single extraction *pulse* — a wild party followed by a bad hangover.

The Lotka-Volterra model, it turns out, is a valuable tool for thinking about this fossil-fuel extraction pulse. That's because, with a little tweak, we can transform the equations into a model of non-renewable resource extraction.

Rather sensibly, the Lotka-Volterra model assumes that the 'prey' population is self renewing — that sheep can replenish their numbers if some get eaten. But this assumption is just a model parameter, and parameters can be changed. If we set the prey replenishment rate to zero, then we transform the 'prey' population into something quite different: it becomes a stock of a *non-renewable resource*.

In this scenario, the Lotka-Volterra model produces a different set of dynamics. Figure 2 illustrates the new pattern. Here, the blue curve represents a non-renewable resource stock. (Let's think of it as a stock of food in a laboratory Petri dish.) And the red curve represents a 'predator' population. (Petri-dish bacteria.) Initially, the bacteria population is small, and so they eat and reproduce merrily. Their numbers swell, and their fixed stock of food declines. When the foodstock reaches a point where further bacteria growth is impossible, the expansion switches to decline. A die off begins, and the bacteria head towards extinction.

The extraction pulse

In Figure 2, the blue curve shows the remaining stock of the non-renewable resource. Now, the model presumes that we know, *in advance*, the total size of this stock. And in the case of a Petri dish filled with food, we certainly do know the size of the recoverable foodstock. However, in more realistic scenarios, the total recoverable resource stock remains uncertain. For example, if the Petri dish is large and the food is distributed unevenly, it may be that the bacteria never reach certain patches. So although this unreachable food 'exists', it doesn't count towards the total stock of recoverable food.

Turning to humans, if the world is a Petri dish, fossil fuels are a 'food' that is unevenly distributed. Sure, we can guess at the total stock of fuel. However, much of this energy will never be extracted, because the cost is too prohibitive. Hence, the total recoverable stock of fossil fuels remains uncertain. For that reason, it's more helpful to look at the dynamics of the Lotka-Volterra model in a slightly different light. Instead of measuring the remaining *stock* of a non-renewable resource, we can look at its *flow* — the rate that it's extracted.

Figure 3 shows this flow-based view. In the case of our bacteria, the flow represents the rate of food consumption. It rises as the bacteria population expands, and then falls as the food gets exhausted and the bacteria population collapses. Now, the point is that in the Lotka-Volterra model, this bell-shaped pulse of extraction is a generic feature of non-renewable resource

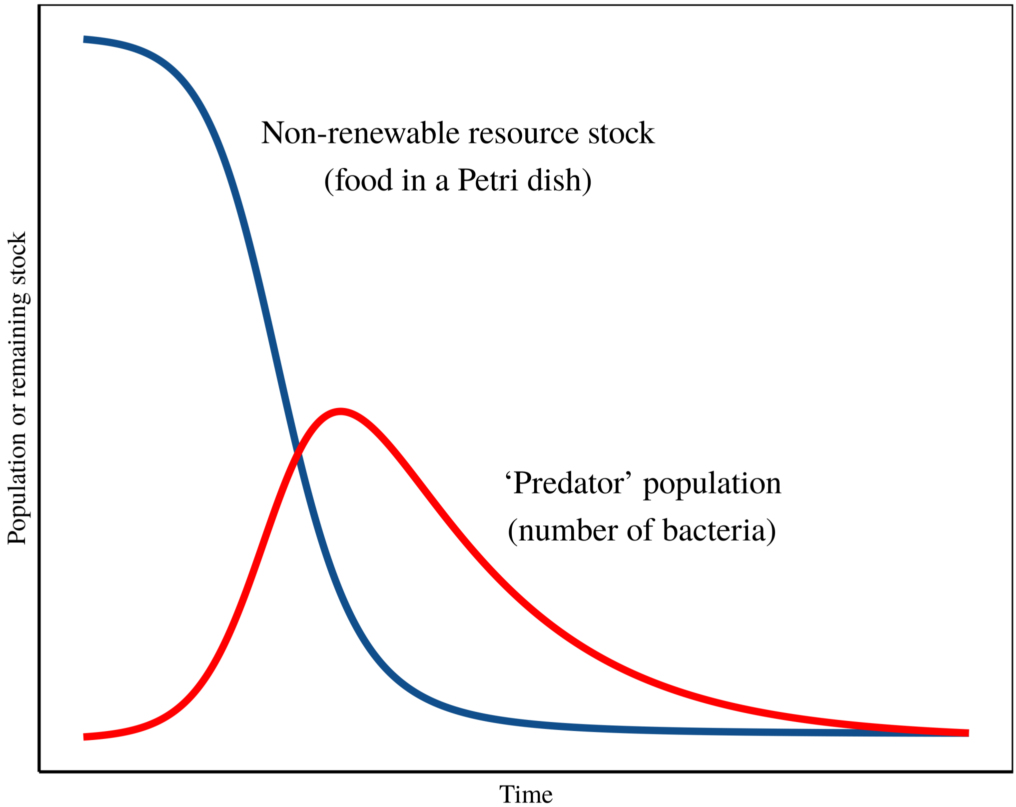

Figure 2: When 'prey' is non-renewable, the Lotka-Volterra model creates a single pulse of predation

Suppose we put a few bacteria in a Petri dish filled with a fixed supply of food. Here's what the Lotka-Volterra model says happens. Initially, the bacteria population (red curve) grows because food is plentiful. But because the foodstock is fixed (blue curve), it gradually declines, eventually reaching a point where further bacterial growth is impossible. At that point, growth switches to decline, and the bacteria population starves to death.

consumption. It applies to humans as much as it applies to bacteria. And unlike the stock-based view (Figure 2), this flow-based view is observable to human participants. We can *watch* the extraction rate of fossil fuels rise and fall. Indeed, in many places, we've already seen both sides of this extraction pulse.

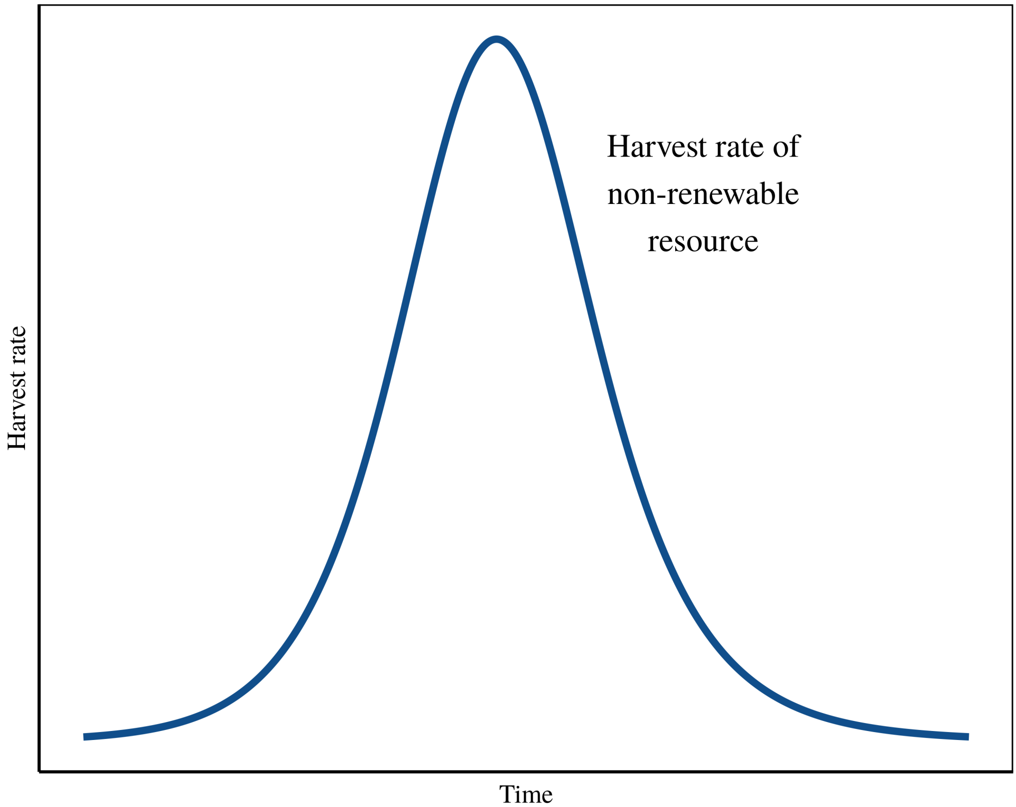

Figure 3: An extraction pulse

Instead of plotting the stock of a remaining resource, this figure plots the *flow* of a non-renewable resource, as predicted by the Lotka-Volterra model. The result is a bell-curved pulse of consumption.

Predatory machines

Continuing the theme of non-renewable resource extraction, if fossil fuels are the 'prey', then who is the 'predator'?

Well, in some sense, the fossil-fuel 'eaters' are literally humans. After all, we grow most of our food with fossil-fuel-based fertilizers, which means that in a way, we 'eat' fossil fuels. That said, the main 'predator' of fossil fuels is not people; it's our fossil-fuel eating machines.

Think about it this way. We use our machinery to wrench fossil fuels from the Earth. Then we feed the harvested energy back to our machines, many of which help us extract more fossil fuels. This loop, it turns out, is exactly the sort of feedback envisioned by the Lotka-Volterra model. To use the Lotka-Volterra equations to simulate the extraction of fossil fuels, we let the 'prey' be fossil fuels; and we let the 'predator' be our extraction technology.

Now, many researchers have realized that fossil fuel extraction can be understood with simple systems models. However, it was Ugo Bardi and Alessandro Lavacchi who first proposed a direct link between the *rate* of resource extraction and the *stock* of extractive technology.

Figure 4 illustrates the connection envisioned by the Lotka-Volterra model. Here, the blue curve plots the extraction rate of a non-renewable resource. And the red curve plots the population of 'predators' — the stock of extraction technology. Notice two things about this simulation. First, both the resource harvest rate and size of the technological stock have a pulse-like behavior — a rise, peak, and fall. Second, the peak of the technological stock *follows* the peak of extraction.

Why this order? According to the Lotka-Volterra model, it's because the extraction technology feeds on the resource being harvested. So when this resource input peaks and declines, the technological 'predators' begin to die off a short while later. (Which is to say that the machines are abandoned and left to rust.)

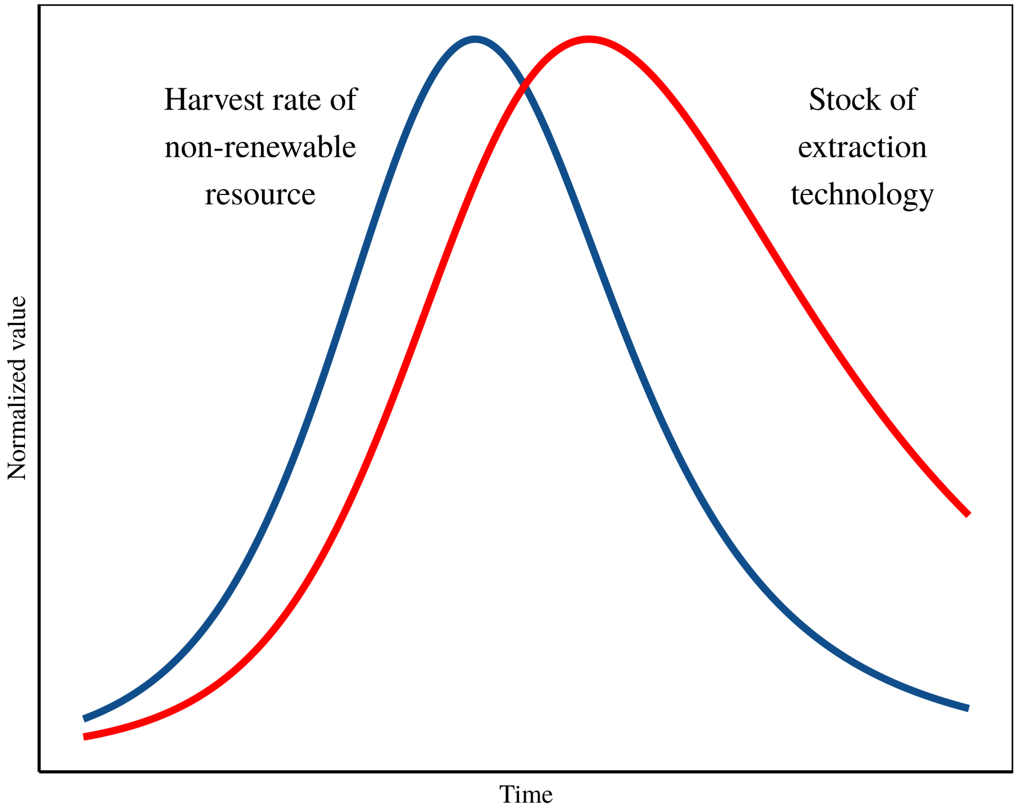

Figure 4: Feeding a technological predator

Here I've plotted the version of the Lotka-Volterra model envisioned by Bardi and Lavacchi. In this simulation, we imagine feedback between the extraction of a non-renewable resource and a stock of extraction technology. In essence, the technology 'feeds' on the resource in question, which means that its fate is linked to the resource itself. The key result is the stock of extraction technology peaks *after* the peak of extraction.

Predators in the Alberta oilpatch

At this point, I'm going to turn to a real-world example of Lotka-Volterra-like behavior. But before doing so, it's worth reminding ourselves that the Lotka-Volterra model is a *toy*. It's purposefully designed to be an over-simplification of reality. So it's somewhat surprising the model has anything to say about the messy arena of human affairs. And yet when it comes to our exploitation of fossil fuels, it seems that humans behave unwittingly like the Lotka-Volterra model predicts.

For a good example of this unintended connection, let's turn to the history of oil-and-gas production in the Canadian province of Alberta. Today, the province is (in)famous for its extraction of unconventional oil from the

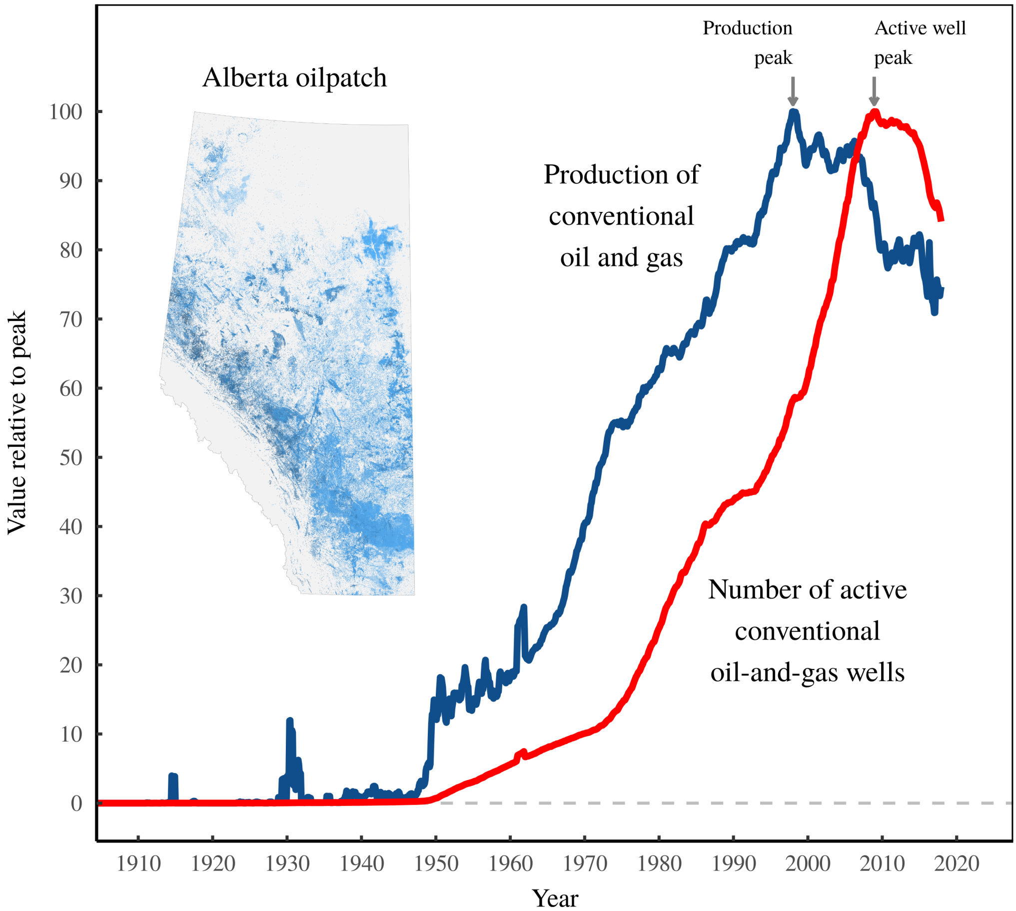

Figure 5: The rise and fall of Alberta convention oil-and-gas production
The blue curve shows the history of conventional oil-and-gas production in Alberta, Canada. The red curve shows the rise and fall of active wells. Like the Lotka-Volterra model predicts, the peak of the technological extractive stock *proceeds* the peak of production. The inset map shows the over 650,000 wells drilled so far. Sources and methods

---

Athabasca tar sands. However, much of the 20th century was spent drilling for conventional oil and gas. Figure 5 shows the history of this geological bonanza.

Following a few false starts in the early 20th century, the Alberta oilpatch got rolling after World War II, driven largely by an unquenchable American thirst for energy. Conventional oil-and-gas production expanded for the next fifty years, but peaked in 1998. Today, Alberta's conventional oilpatch is in steep decline, and the big players have largely moved north to the unconventional oil sands.

Now, the oilpatch is driven by a simple principle, which is that you extract oil by drilling holes in the ground. The more holes you drill, the more oil you get. Or at least that's how it works at first. Over time, the big reserves get depleted, and more and more wells become duds. Eventually, there are enough duds that oil production begins to decline even though the number of wells continues to increase. When that happens, the economics of the oilpatch shift. Drilling slows, unproductive wells are left to rust, and the number of active wells begins to decline.

The red curve in Figure 5 shows this pattern of active-well peak and decline. It is eerily similar to the Lotka-Volterra model in Figure 4. The message here is that the players in the Alberta oilpatch seem to be unwitting puppets of a toy model. As predicted by the Lotka-Volterra model, the stock of Alberta's active oil-and-gas wells peaked shortly *after* the peak of oil-and-gas extraction.

Now, the effect of a good chart is always to collapse complicated behavior into a graphical pattern that's simple enough to comprehend. So when we stare at a chart like Figure 5, it's easy to lose sight of the antics beneath it. For that reason, I've included a map of the Alberta oilpatch, where each oil-and-gas well is an imperceptibly small dot. Today, there are over 650 000 wells in total, each with its own story of ambition, glory, and failure. Importantly, there was no grand plan to the Alberta oilpatch, other than to make money selling the riches of the Earth. But ironically, it's this *lack* of plan that gives rise to the overarching pattern of rise and fall.

The Lotka-Volterra model assumes a basic instinct to eat when the pickings are good, and starve when the food runs dry. But it could be that humans, in all our wisdom, are able to suppress this urge. For example, we can imagine a scenario in which the government sets quotas on oil-and-gas drilling — quotas designed to keep production constant. In the face of such planning, the Lotka-Volterra model would have nothing to say about oil extraction.

Although humans are surely smart enough to enact such policies, rarely do we actually *do* so. Instead, when faced with a stock of exploitable resources, we're gripped by an animalistic urge to consume them as fast as possible. The Lotka-Volterra model captures this urge, which is why it seems to predict the large-scale pattern of how we extract resources, without knowing anything about our small-scale antics.

Shocking the system

When we 'play' with a model, it's important to be open about its limitations. On that front, the Lotka-Volterra model is an obvious over-simplification of the real-world, which means we expect to find many situations where it breaks down.

For example, we can imagine a population of sheep and wolves in which a farmer drastically culls the wolf pack. Or we can imagine a bacteria-filled Petri dish in which a scientist suddenly dumps in more food. Neither situation can be anticipated by the Lotka-Volterra model, which pretends that its modeled populations exist in splendid isolation. In the arcane language of economics, these system shocks are said to be 'exogenous'; they are not part of the Lotka-Volterra model, which means they can't be predicted.

That said, these shocks can be put in 'by hand'. To add a system shock to the Lotka-Volterra model, we can arbitrarily change the predator/prey population midway through the model run. Then we see how the model responds.

To get started with system shocks, let's return to our example of Petri-dish bacteria which are busy eating a finite stock of food. Left alone, the bacteria's food consumption will follow the familiar resource 'pulse', plotted in Figure 3. Food consumption will rise as the bacteria population expands, and then collapse as the bacteria starve. Now suppose that partway through this consumption pulse, a benevolent scientist dumps more food into the dish. What happens?

Well, it turns out that the impact depends on the timing of the food dump. If the food dump happens *early* in the experiment, the shape of the consumption pulse remains essentially unchanged. Figure 6 illustrates. Here, the foodstock quadruples early on, before the bacteria population has had much time to grow. The result is a minor uptick in food consumption, followed by the expected pulse of growth and decline.

In contrast, if the food dump happens *late* in the experiment, the effect is drastically different. Figure 7 illustrates. Here, our scientist waits until the original foodstock has begun to wane before dumping in a new bonanza. The result is a massive increase in resource consumption, which creates a second extraction pulse.

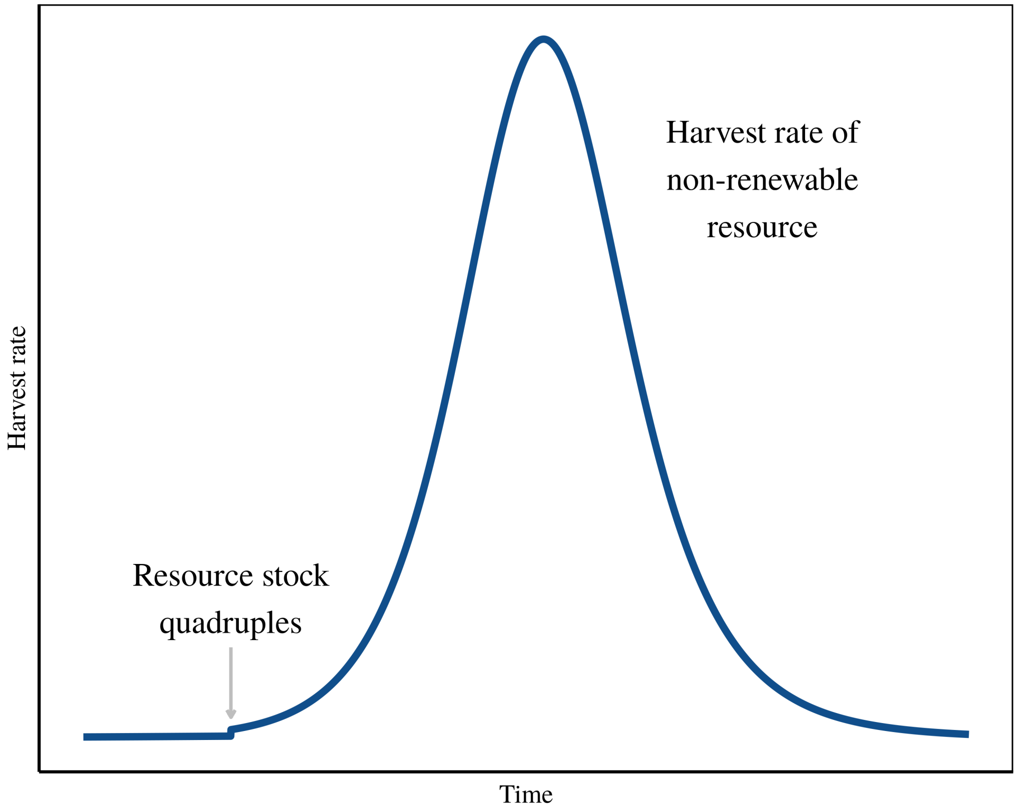

Figure 6: An early-game resource shock

Here we imagine bacteria in a Petri dish with a finite stock of food. Early in the consumption pulse, a scientist quadruples the foodstock. According to the Lotka-Volterra model, not much happens. That's because at the time of this early dump, the bacteria population is small, so it can't do much with the extra food. So the consumption pulse plays out as though the larger foodstock was there all along.

So why does it matter *when* we dump in the food? Well, because the bacteria's ability to harvest food depends on their *population*. If we add food when the population is small, the bacteria can't do much with it — their numbers are too few. If, however, we dump food into the dish late in the game, there is a large population of starving bacteria ready to gobble up the resource.

Returning to humans, the same scale principle holds. For example, suppose that in 1870, a benevolent god somehow quadrupled the global stock of oil. Would anyone have noticed? Probably not. At the time, oil extraction was in its infancy; our seismic technology was non-existent, our drilling technology was juvenile, and our refining technology was rickety. In short, when faced with an early-game quadrupling of our oil reserves, pretty much nothing would have happened (at the time).

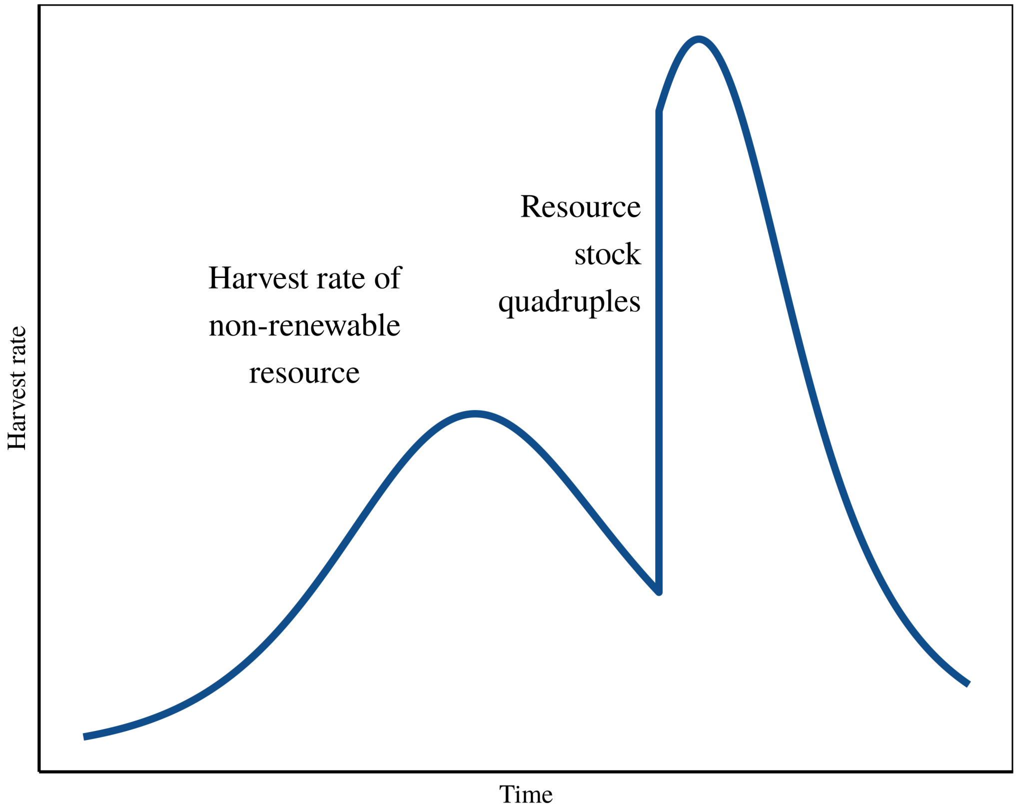

Figure 7: A late-game resource shock

Unlike an early-game resource shock, a late-game shock changes the shape of the consumption pulse by adding a second peak. Here, we imagine a population of Petri-dish bacteria left alone to eat a finite stock of food. After food consumption has peaked, a benevolent scientist quadruples the remaining foodstock. This bonanza creates a second peak of consumption — one which burns more brightly and more briefly than the first peak.

Now imagine that the remaining stock of oil quadrupled *today*. Would anyone notice? It's a silly question that's not rhetorical. As it happens, the United States is in the midst of an oil-and-gas bonanza — one created by the exploitation of tight oil and shale gas. Of course, these reserves haven't just popped into existence — they've been there all along. What changed is our *technology*.

For most of the 20th century, oil and gas was extracted by drilling a vertical well, and then sucking out the reserve. This technique works well if the formation is porous enough for the oil to flow. But if the formation is imporous, the well will come up dry. Now suppose that instead of drilling vertically, we bent the borehole and extended it horizontally though the reserve. And

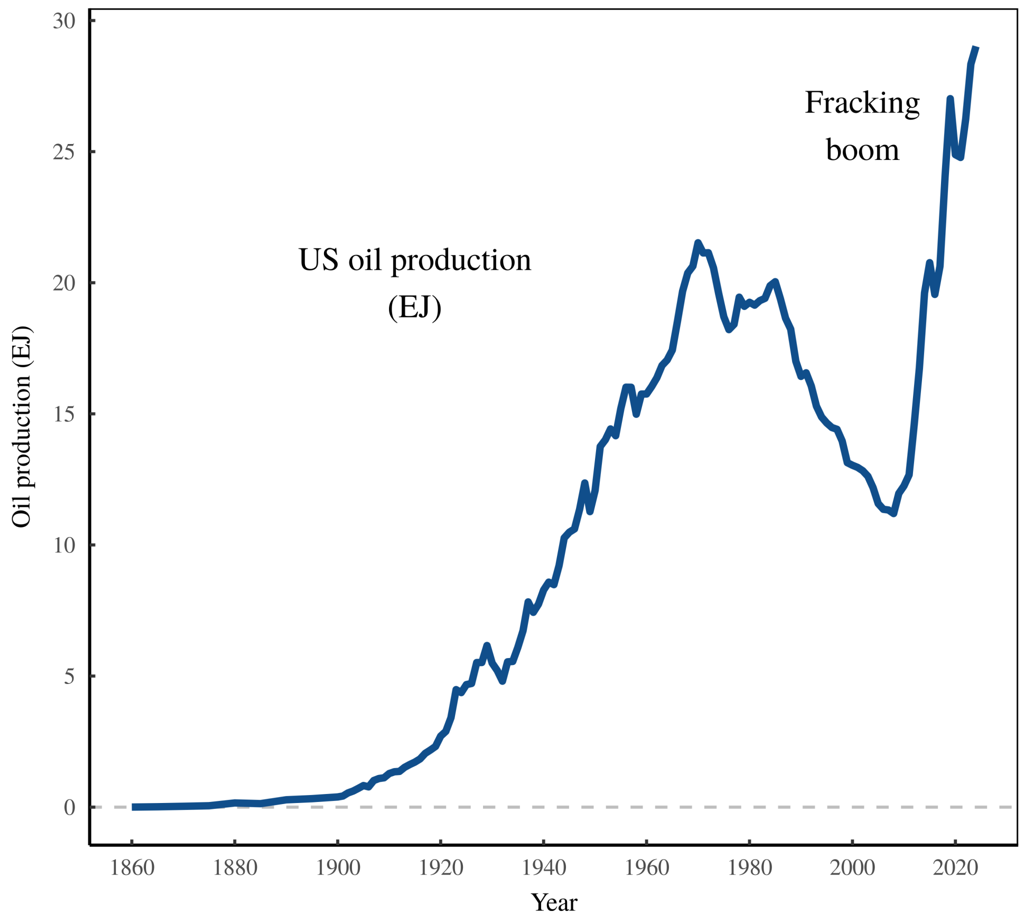

Figure 8: A second bonanza — oil production in the United States

After more than three decades of decline, US oil production exploded in the 2010s. The turnaround owes to the new technique of fracking — using high-pressure liquid to fracture oil formations that were previously too impervious to flow. Sources and methods

---

then suppose that we pumped in a high-pressure liquid which fractured the formation. Well then, this previously inaccessible resource would flow like melted butter. That, in a nutshell, is how the fracking revolution has worked.

To see this revolution, let's turn to Figure 8, which plots the history of US oil production. For decades, the United States was the poster child for peak oil. In 1956, the geologist M. King Hubbert predicted that US oil production would peak in the early 1970s. And that is exactly what happened. From the next four decades, production declined. But in the mid 2000s, the fracking revolution opened up new reserves, sending US oil production to new heights.

Now the pertinent question is — how long will this second oil-and-gas bonanza last? For their part, peak oil theorists have become less strident than they were in the mid 2000s, in large part because the fracking revolution has shown the importance of technological change. That is, the amount of recoverable oil is affected not just by the Earth's geology, but also by the tools we use to harvest fossil fuels.

So while I won't give a definite prediction for the second peak of US oil production (I've already made one that's proved wrong), the Lotka-Volterra model does give us reason to be bearish about the timing of this peak. Looking at Figure 7, the Lotka-Volterra model predicts that a late-game resource shock creates a second consumption peak that burns more brightly and briefly than the first peak. Again, this is because when resources are added late in the game, there's a huge stock of 'predators' ready to exploit them.

Likewise, in the United States, the fracking revolution is taking advantage of an immense technological stack that is hungry for oil (and that had been slowly starving for decades). Because of this latent capacity, the US will likely burn through its unconventional oil reserves more quickly than it did with the conventional stuff.

Killing off predators

Continuing the theme of system shocks, so far we've explored what happens if we shock the Lotka-Volterra model by adding new 'prey'. In the same vein, we can also shock the model by killing off 'predators'.

As before, what interests me is how this shock plays out in the context of a non-renewable resource 'pulse'. Since it's the 'predators' that do the consuming, killing them off has the predictable effect of temporarily dampening resource consumption.

Figure 9 illustrates the effects of a predator die off. Returning to our example of Petri-dish bacteria, suppose that midway through their growth pulse, a vindictive scientist dumps cyanide into the dish. The bacteria population crashes, as does their rate of food consumption. However, the population soon recovers, and the consumption pulse begins again.

Similarly, anything that kills off humans — be it famine, pandemic, or war — will disrupt our consumption of resources. And the disruption will be doubly severe if it also destroys our technology, as is the case during warfare. In this regard, fossil fuels have been both a blessing and a curse. On the one hand,

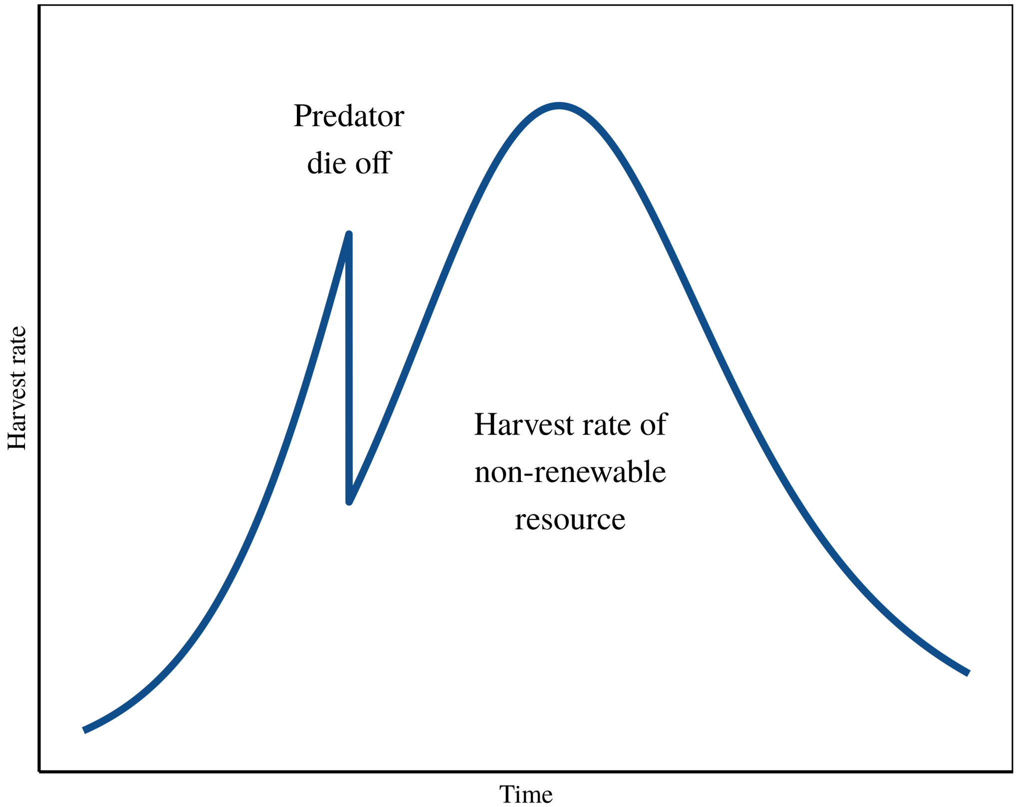

Figure 9: Killing off predators

Here we suppose that early in the consumption pulse of a non-renewable resource, most of the predators die off. The resource harvest rate plunges, but not for long. As the predator population recovers, the pulse of resource consumption resumes its course.

---

they've enabled an unprecedented rise in our standard of living. But on the other hand, they've magnified our destructive power, making war far more terrifying.

Perhaps more than any country, Japan offers the best example of how a consumption pulse can go wrong when it is interrupted by war. For much of the early 20th century, Japan was busy doing what every other great power had done before, which was to conquer new territory. Japan's main sin was that it was late to the imperial game, which meant its expansion tread on the toes of established colonial powers.

As part of this imperial game, the Japanese military decided, in late 1941, to prod the US empire by bombing Pearl Harbor. It was a foolish decision. At the time, the United States was by far the world's most powerful country, consuming about a third of the world's energy. So egging it into war was destined to end badly.

And for Japan, end badly the war did. Not only was much of the country flattened by conventional bombs, Japan remains the only population to have been devastated by a nuclear bomb. Figure 10 plots the scale of this wartime destruction, measured in terms of Japan's share of world energy use. After prompting the US into World War II, Japan's share of world energy use plummeted. It didn't recover to pre-war levels until 1966.

Now, I realize that it feels crass to reduce the violence of war to an abstract systems model. But really, it's no more crass than converting the violence of predation into a mathematical equation — which is exactly what the Lotka-Volterra model does. In this case, the destruction of imperial Japan offers a case study in what happens when machines and infrastructure (our technological 'predators') get destroyed.

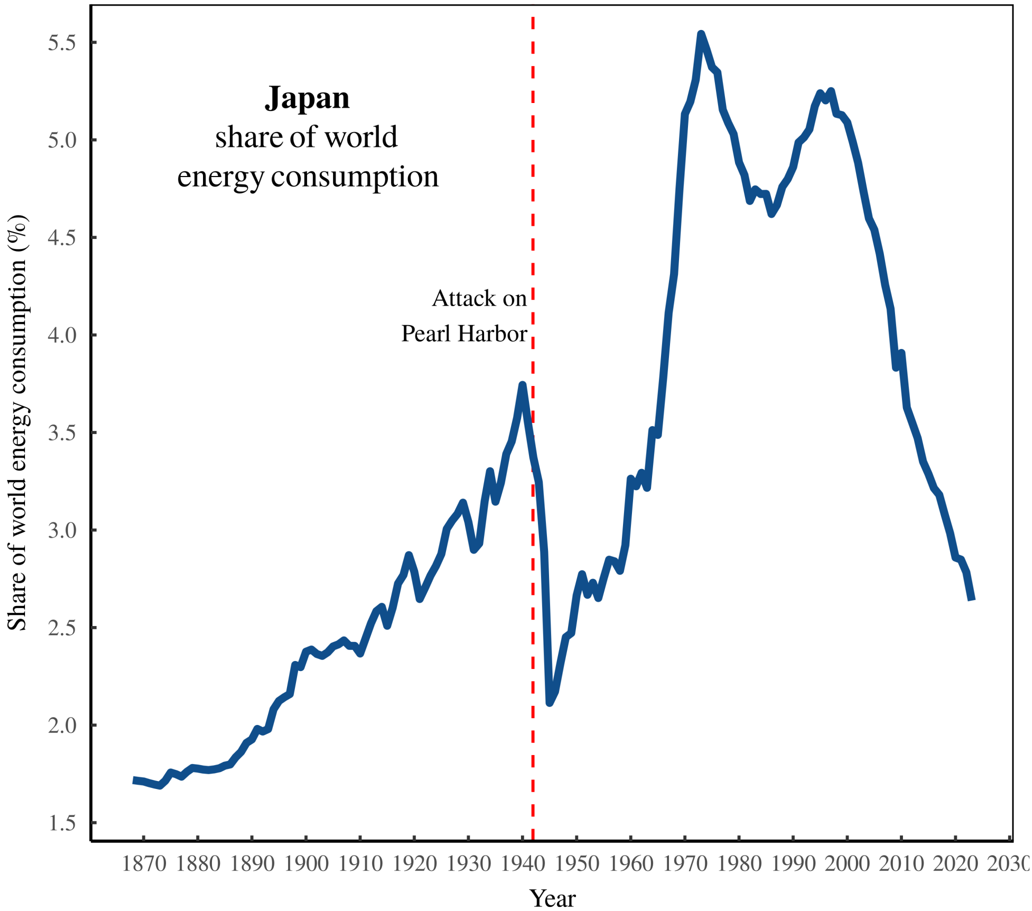

Figure 10: Japan's share of world energy use and the devastation of WWII

It's easy to spot the moment when Japan provoked the US into declaring war (in late 1941). The ensuing US bombardment decimated Japan's infrastructure (and of course, its population), sending Japan's share of world energy use back to 1890 levels. The post-war recovery took decades. Sources and methods

Social collapse

Since human-built machines don't (yet) have a life of their own, they needn't be destroyed to be rendered useless. Anything that obstructs their human caretakers will have the same effect, freeing technology to sit idle. For this reason, periods of social collapse can (like war) be treated as a 'predator shock' — moments when our active technological stock suddenly decreases.

For example, during the Great Depression, much of the world's machinery lay unused, simply because people couldn't afford to use it. More recently, the collapse of the Soviet Union provided a similar experiment with idle technology driven by social chaos.

When the Soviet regime died in 1990, Western economists confidently declared that markets would emerge and pick up the slack left by the absence of state planning. Unsurprisingly, the market 'miracle' worked rather differently. In the aftermath of the Soviet collapse, former member states suffered a severe and prolonged depression.[2]

The history of Russian oil production offers a good window into this collapse. During the later years of Soviet control, Russian oil production had exploded, reaching a pinnacle in 1987. But in the aftermath of the Soviet collapse, Russian oil production was cut nearly in half. It didn't return to the Soviet-era high until 2019. Figure 11 illustrates this free market 'miracle'.

Again, I think the Lotka-Volterra model provides some useful insight into the post-Soviet depression. Sure, it says nothing about how or why the Soviet Union collapsed. But when the ensuing social chaos rendered a large portion of Soviet machinery inoperable, we can treat the catastrophe that followed like a sort of 'predator' die off. Only recently, have Russia's 'oil predators' recovered.

[2] Dmitry Orlov's book *Reinventing Collapse* gives a fascinating account of the suffocating environment in post-Soviet Russia.

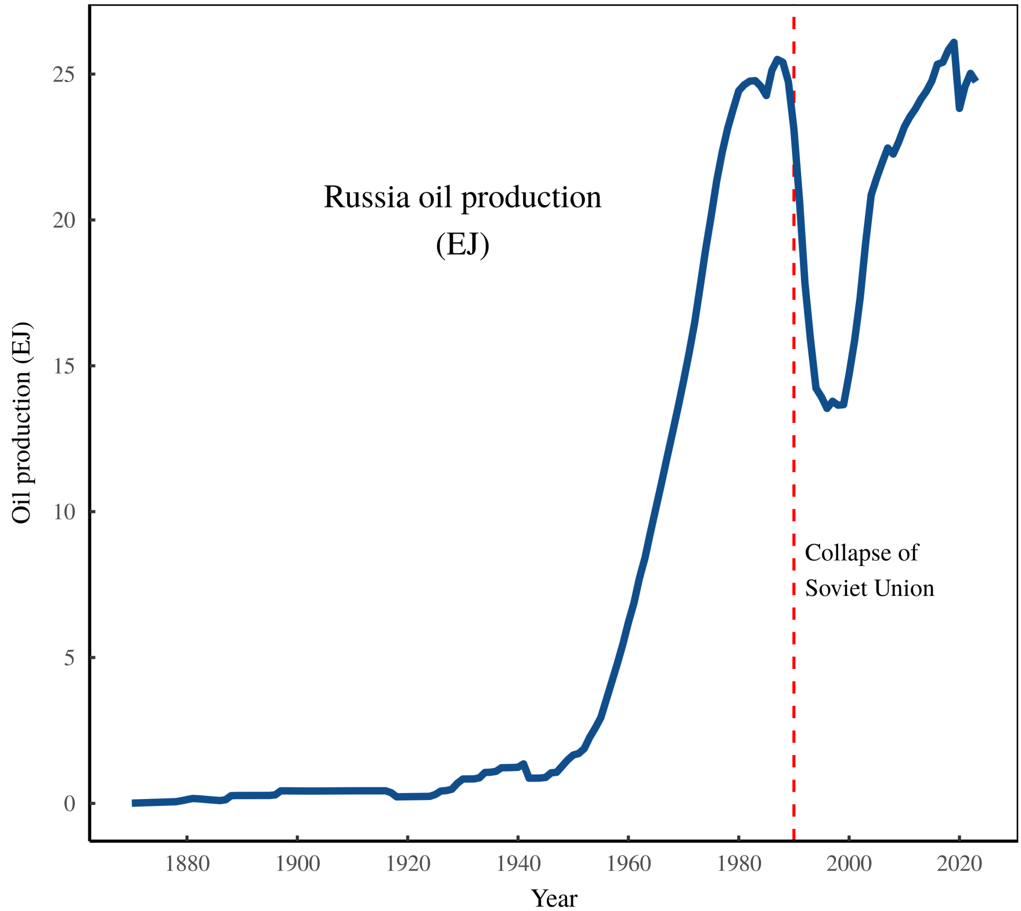

Figure 11: A market 'miracle' — Russian oil production implodes following the Soviet collapse

When the Soviet Union dissolved, Soviet-block countries experienced a severe and prolonged depression, visible clearly in the production of oil. For example, Russia's oil production was cut nearly in half, and didn't return to the Soviet-era peak until 2018. Sources and methods

A humble toy

This concludes my tour of the Lotka-Volterra model, which I'll remind you, is best considered a humble mathematical toy. The Lotka-Volterra model doesn't make grand predictions about the future. It's not compelling enough to attract acolytes. It's not seductive enough to be enshrined in official dogma. And it's not enthralling enough to be the subject of political debate. No, the Lotka-Volterra model is a simple thought experiment about the effects of population feedbacks.

And yet, I hope to have convinced you that the Lotka-Volterra model is *useful*. Life on Earth is dominated by feedback effects, and it behooves us to understand how they work. The Lotka-Volterra model offers a comprehensible entry point into the world of systems modeling, a world in which simple principles generate complex effects. In an academic landscape dominated by neoclassical economic fantasies, surely we could use more of this type of thinking.

Support this blog

Hi folks. I'm a crowdfunded scientist who shares all of his (painstaking) research for free. If you think my work has value, consider becoming a supporter.

Become an ETD supporter

Minsky

It would be foolish to write a post about feedback modeling without mentioning the hard work being done by heterodox economist Steve Keen.

Backing up a bit, systems models consist of nothing but sets of differential equations. In the early days of modeling, scientists coded these equations by hand. But that gets tedious quickly. And so researchers developed graphical tools for creating systems models. Today, there are many such tools, but virtually all of them are proprietary and closed source. The exception is the systems modelling program *Minsky* developed by Steve Keen and coded by Russell Standish. *Minsky* is free and open source, and designed specifically with economics in mind. I encourage you to try it.

The Lotka-Volterra equations

The Lotka-Volterra model consists of a set of coupled differential equations, typically written as:

$$\frac{dx}{dt} = \alpha x - \beta x y$$

$$\frac{dy}{dt} = -\gamma y + \delta x y$$

Here, x is the prey population and y is the predator population, while dx/dt is the rate of change of the prey population and dy/dt is the rate of change of the predator population. The remaining terms are model parameters which, to me, make more sense if we rewrite the model in terms of population growth rates.

To reframe the Lotka-Volterra equations in terms of growth rates, note that a growth rate is simply a rate of change expressed as a portion of the thing changing. So the growth rate of x is:

$$\widehat{x} = \frac{dx/dt}{x}$$

Likewise, the growth rate of y is:

$$\widehat{y} = \frac{dy/dt}{y}$$

When we reframe the Lotka-Volterra equations in terms of growth rates, they simplify as follows:

$$\widehat{x} = \alpha - \beta y$$

$$\widehat{y} = -\gamma + \delta x$$

In English, these equations state:

1. Without predators, the prey population will grow exponentially at rate $\widehat{x} = \alpha$.

2. Without prey, the predator population will decline exponentially at rate $\widehat{y} = -\gamma$.

3. The presence of predators y decreases the growth rate of prey by $-\beta y$.

4. The presence of prey x increases the growth rate of predators by δy.

And that's it! From these simple equations comes a host of dynamics that are not predictable from algebra alone.

Sources and methods

Alberta oilpatch (Figure 5)

This data is from my post 'A Case Study of Fossil-Fuel Depletion'. Detailed methods are available here.

US oil production (Figure 8)

Oil production data is from the following sources:

- 1949 to present: Energy Information Agency, Table 1.2, Primary energy production by source
- 1860 to 1948: Historical Statistics of the United States, Table DB157

Japan's share of world energy use (Figure 10)

World energy use is from Our World in Data, Energy Production and Consumption.

Japan's energy use is from the following sources:

- 1965 to present: the Energy Institute's Statistical Review of World Energy, 2025
- 1900 to 1964: Benjamin Warr's REXS dataset
- 1868 to 1899: estimated from carbon emissions per capita. Both the emissions data and the population data is from Our World in Data

All Japanese series are indexed backwards from the Statistical Review data. Also note that the carbon-based estimates are indexed twice. First, I index the carbon data to Japanese energy use in 1900. The resulting series assumes that Japan's energy use prior to 1900 directly tracks its carbon emissions. The problem with this estimate is that it ignores non-fossil fuel sources of energy, which become more important as we head back in time. To correct this problem, I then add to the time series the constant value of 26,000 KCal of energy per person per day. Finally, I re-index this updated energy estimate to the statistical data from 1900.

(For those who are interested, I've used the same carbon-based method to estimate the historical energy use of the Soviet Union.)

Russian oil production (Figure 11)

Russian oil production data is from the following sources:

- 1985 to present: the Energy Institute's Statistical Review of World Energy, 2025
- 1870 to 1984: data digitized from Figure 1 in John Grace's book *Russian Oil Supply: Performance and Prospects*

Further reading

Bardi, U., & Lavacchi, A. (2009). A simple interpretation of Hubbert's model of resource exploitation. *Energies*, *2*(3), 646–661.

Orlov, D. (2008). *Reinventing collapse: The Soviet example and American prospects*. New Society Pub.