Unit 6: Sociopolitical cycles

READINGS AND RESOURCES

- Wangersky PJ (1978) Lotka-Volterra population models. *Annual Review of Ecology and Systematics* 9(1): 189-218.
- Turchin P (2003) The metaethnic frontier theory. In *Historical dynamics* (pp 50–77). Princeton University Press.

Much of life is cyclical

We've modeled a number of dynamical social systems in this course. However, one thing pretty much all the systems we've looked at have had in common is the tendency to evolve toward a *stable equilibrium*. Yet social life is pretty much never at equilibrium, despite what your friendly neighborhood economist might tell you. Instead, much of life is cyclical.

- A time of war, a time of peace
- A time to dance, a time to mourn
- A time to cast away stones, a time to gather stones together (etc)

In this last modeling unit, we're going to explore a couple of models that allow for cyclical behavior. What is needed for cyclical behavior is *feedback* and *delay*. That is, the dynamics that give rise to growth must not only contribute to eventual decline, but must also do so with enough of a lag that growth does not simply reach a stable equilibrium, but exceeds the conditions for stability and must therefore decline, with decline also happening too rapidly to slow down in time for equilibrium.

There are a number of mechanisms that can give rise to cycles. The most well known is the Lotka-Volterra predator-prey model developed in the 1920s, which I discussed briefly in Unit 1. The purpose of this model was to explain the widely observed phenomenon among hunters and fur traders that populations of predator and prey species tended to repeatedly rise and fall on similar time scales. For example, Hudson's Bay Company in Canada noticed cyclical patterns in the populations of the animals they hunted, like the Canada lynx and the snowshoe hare. Not only did both populations rise and fall, but their cycles seems to be locked in phase, and unexplained by external factors like climate. The model is based on coupled differential equations, but its assumptions are very simple: In the absence of lynx, hare will reproduce and increase in number. In the absence of hare, lynx will starve and decrease in number. When both are present, lynx will eat hare, increasing their numbers while diminishing the hare - their food source.

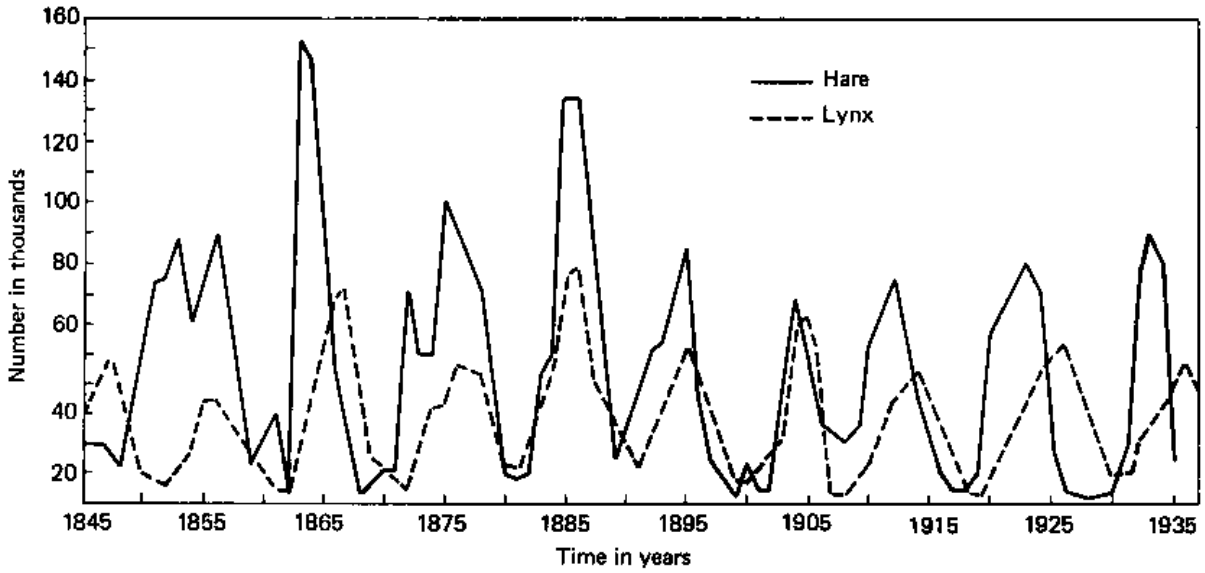

Figure 9-3. Changes in the abundance of the lynx and the snowshoe hare, as indicated by the number of pelts received by the Hudson's Bay Company. This is a classic case of cyclic oscillation in population density. (Redrawn from MacLulich 1937.)

Here, we will begin by exploring a related model: the ecological host-pathogen model. This is a very simple model, but can lead to complex and counterintuitive outcomes. After exploring this simple model, we will explore one that is considerably more complicated but operates of similar principles: that of the rise and fall of empires.

The host-pathogen model

CODE: hostpathogen.nlogo

In a sense, we're going back to the first system we modeled: the spread of contagion. Then, we considered the time course of a contagion spreading through a population. We even considered that the host could recover and become uninfected. But we always assumed a fixed population of agents that neither died nor reproduced. And that's fine for some problems. However, it's now time to think a bit more ecologically.

In this model, there is a host species that can become infected by a lethal pathogen. Spatial locations on a grid can represent either susceptible hosts, infected hosts, or empty locations onto which hosts can reproduce. Pathogens infect hosts. The pathogen spreads by passing from one host to another, and an infected host will eventually die from illness. If there are no nearby hosts, the pathogen cannot spread. Hosts reproduce when they are healthy, but not when they are infected.

Our model assumptions lend themselves to other possible interpretations. As noted, the model is similar to the Lotka-Volterra predator-prey model, and this similarity can be embraced more directly, with the hosts corresponding to prey and the pathogens to predators. In this interpretation, prey reproduce in the absence of predators but are killed by predators, while predators reproduce after killing prey as long as another prey population exists nearby, but they die in the absence of prey. Another interpretation is possible in terms of human populations: consumers and resources. Consumers use up renewable resources to reproduce, which depletes the resources. In the absence of resources, consumers perish, but resources renew in the absence of consumers. For simplicity, however, I'll continue to talk about this model in terms of hosts and pathogens.

This is actually one of the simplest models to code. It's so simple you may ask why I didn't *start* the course with this model. One answer is that it allows me to point out that more complicated models aren't always better or more sophisticated. You are not necessarily a better modeler if your models are more complex. Rather, strive for the simplest possible model to capture something interesting about your system. Another reason is that although the model assumptions are simple, its dynamics can appear quite complex. A final reason is that the final model we will consider in this unit IS very complicated, and the host-pathogen model provides a useful bridge to it.

For this model, we'll use a square grid as we've been doing, with agents paying attention to their nearest four neighbors (up, down, left, and right). The cells of the grid will be the agents, which can take on one of three states: empty, susceptible, and infected. The model is pretty much fully defined by the rules governing the transitions from one state to another, so let's take a closer look at how those transitions are formalized.

Transition rules:

- Empty cells can become susceptible hosts. An empty cell can become a susceptible host if at least one of its neighbors is one, and so reproduces onto it. A susceptible host attempts reproduces onto a neighboring empty cell with probability r, the host *reproductive rate* (we say "attempts" because at maximum only one neighbor can be successful). In our model, all hosts are identical, so the probability of going from an empty cell to a susceptible host increases with the number of neighbors that are susceptible. The probability of transitioning is equal to one minus the joint probability that none of one's susceptible neighbors attempts to reproduce. If n_S is the number of susceptible neighbors (always between 0 and 4 in our model), the probability of transitioning from an empty cell to a susceptible host is: $P(0 \to S) = 1 - (1-r)^{n_S}$.
- Susceptible hosts can become infected hosts. A susceptible host becomes infected by one of its infected neighbors with probability τ ,

the *transmissibility* of the pathogen. As before, the probability of becoming infected increases with the number of infected neighbors, and is equal to one minus the joint probability that none of one's infected neighbors transmits the pathogen. If n_I is the number of infected neighbors, the probability of transitioning from a susceptible host to an infected host is: $P(S \rightarrow I) = 1 - (1 - \tau)^{n_I}$.

- Infected hosts can become empty cells. Infected hosts die with a probability equal to *v*, the *virulence* of the pathogen. This transition probability does not depend on the state of one's neighbors, only on being infected.

Our model therefore has three parameters governing how the cells transition from one state to another. In our version of the model, empty cells will be black, susceptible hosts will be green, and infected hosts will be red. We'll initialize our model with a healthy population of susceptible hosts, when an infection breaks out in the middle.

SETTING UP THE MODEL

- *reproductive-rate* slider
- *transmissibility* slider
- *virulence* slider

INITIALIZATION

- Set all patches to green (susceptible)
- Infect a single patch in the center (red)

DYNAMICS

- ASK PATCHES:
  - Consider my current state and the state of my neighbors, and transition with appropriate probability as defined above.

PLOTTING:

- We can see agent dynamics on the grid as cells transition between colored states. We'll also plot the populations of susceptible and infected hosts over time.

RESULTS

- For a very wide range of parameter values, you can get neat spatial dynamic patterns. I encourage you to play around with this model. It's fun to watch.
- In some cases we can get rather substantial oscillations in the host and pathogen populations. Use ($r = .01, \tau = .15, v = .1$). So what happens? The pathogen spreads from host to host. However, the virulent pathogen kills its hosts rapidly, limiting the number of new hosts available to infect, and fizzling out in local populations. This allows

susceptible hosts to recover, providing new fertile soil for the spread of infection. This manifests as cycles widespread infection and recovery.

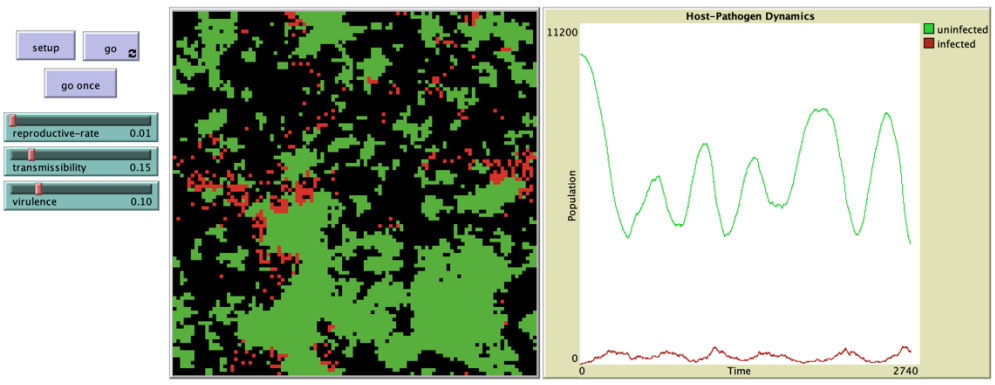

- Exploration of the model parameters provides some interesting and possibly counterintuitive results. For example, you might think that an increased rate of reproduction for susceptible hosts would provide for a robust population and diminished rates of infection. However, the opposite is the case here. Higher reproductive rates for hosts actually favor the pathogen. To see this, increase the reproductive rate from $r =$ .01 to $r =$ .1). This is because there are now more susceptible hosts to which the pathogen can spread, without empty regions persisting long enough to block pathogens from spreading before they kill their hosts. The spatial display shifts from the army of pathogens moving methodically over the landscape to a fierce boil of life and death. In this figure, I've increased r as the model runs.

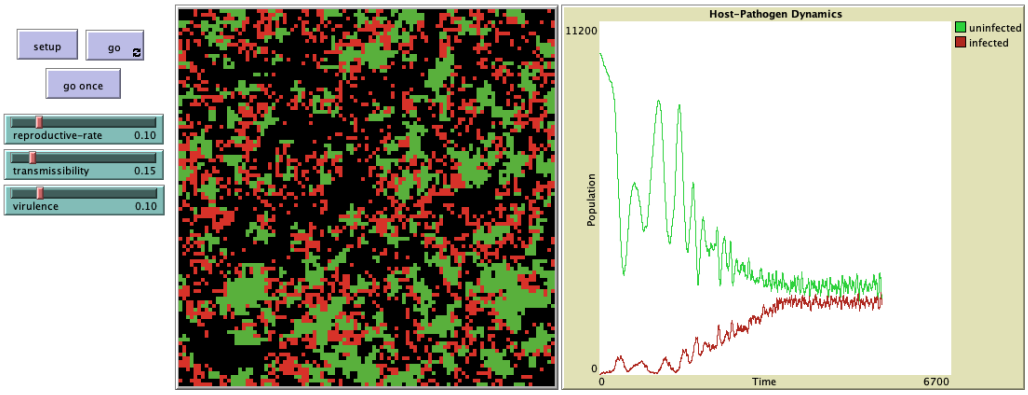

- So the rate of infection has increased with higher host reproduction. What we probably want is a less virulent pathogen so fewer hosts are killed, right? Again, the opposite is the case. Lower the virulence from $v =$ .1 to $v =$ .04. Infected hosts now outnumber uninfected hosts, and the number of empty cells has boomed. Instead of the infection spreading through a health population, it looks almost as if the susceptible hosts are fleeing the rampaging infection. This happens

because infected hosts don't die off as quickly, but stick around to infect more of their neighbors, enhancing the spread of the pathogen. In the figure below, I've continued the run above and decreased the virulence. If the virulence is lowered still more to $v = .01$, the population reliably collapses.

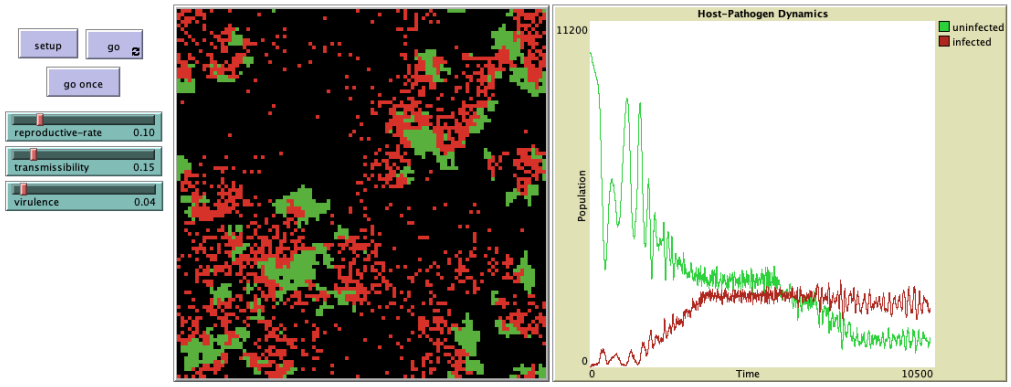

In general, this model has very rich dynamics for something so simple. I encourage you to play around with it. Nevertheless, it might not be clear how what we've done applies to the sorts of things social scientists who are not epidemiologists are interested in. (I think it's interesting, but your mileage may vary). So, as our final model in this course, let's consider something a bit more ambitious: the rise and fall of empires.

The metaethnic frontier theory

Much of the history of civilization—the last 5 to 10 thousand years of human history—can be viewed in terms of the rise and fall of empires. The general story that is often told is one of continued growth. After the agricultural revolution ~11,000 years ago, humans in many parts of the world became more tied to location and saw rapid increases in their population size. This newly-increased population facilitated division of labor and organizational challenges that were in turn well solved by larger and more hierarchical governance structures. However, people do not so readily give up their group identities and join with other groups. Moreover, however, with growth also comes new types of scarcity, and that breeds competition. The result is war, which is widely thought to be a major driver of increased societal complexity. The strong version of this claim can be summarized by the sociologist and historian Charles Tilly: "War made the state and the state made war." With war comes conquest, and therefore we see not only the rise but also the dissolution of empires. But why exactly? Why would competition give rise to cycles of empires rather than more stable equilibria? We need to turn our ideas into models.

How do we model this sort of thing? History definitely seems far messier than most if not all of the systems we've tackled so far. The short answer is that it's very difficult. But difficult is not the same as impossible. If historical dynamics are subject to patterns and forces that recur with at least a modicum of regularity, then those dynamics are modellable. As always, we can start with a model that is drastically simple but nevertheless captures some of our key assumptions about how our system works.

In his 2003 book, *Historical Dynamics,* the ecologist-turned-historian Peter Turchin attempted to explain some of these cyclical dynamics using what he called the metaethnic frontier model, which develops the idea that empires foster a superordinate or metaethnic identity among their citizens. The theory is actually an extension of one developed by the 14th century Arab scholar Ibn Khaldun. A key concept in the theory is *asabiya*, an Arabic word denoting a sort of social solidarity, with an emphasis on unity, group consciousness, and sense of shared purpose. Note that asabiya is a property of a population, not an individual. It is therefore a *group-level trait*, which I discussed at length in Smaldino (2014). Group-level traits are inherently challenging to model as emergent phenomena. So, like many other models, we will just take the trait as given.

The theory proposes that polities that share borders with different polities will be bolstered in their sense of asabiya, which will be important for producing group-level phenomena like warfare and empire building. As empires expand their territories, their increased access to resources will drive further growth. However, as empires become larger, asabiya among those individuals who live far from the frontier will decline as a result of increased internal competition and a lack of exposure to external threats. Meanwhile, the constraints of expanded infrastructure mean that maintaining a frontier becomes more difficult as the empire's size expands. Eventually, an empire can be stretched too thin, and becomes susceptible to conquest.

The theory says nothing about how the first empires form, nor does it account for other factors, like geography, topography, and technology. Nevertheless, it is grounded in pretty well-founded notions of human society, and is clear enough to be formalized. Formalization will help us see how precise and plausible the theory is.

The metaethnic frontier model

CODE: metaethnic.nlogo

The model works as follows. The smallest spatial unit in the model, a cell on a square grid, corresponds to a small regional polity or chiefdom. However, these polities can also become absorbed into a large territorial empire. Cells

have an imperial index denoting the empire to which they belong; independent chiefdoms have an index of 0.

Each cell has a level of asabiya, *S,* which represents the average degree of collective solidarity felt by its population, and is determined by the position of the cell relative to its imperial boundaries. Asabiya grows logistically to an upper limit of 1 if it is situated next to a boundary between two empires, or between an empire and a chiefdom. If the cell is not on a frontier, asabiya declines exponentially to 0. In determination of frontier status, the cell considers only its four nearest spatial neighbors.

Territorial empires are characterized by two numbers: the number of regions they control (Ai), and the average asabiya of those cells, S-bar. At each time step, each cell considers an attack on its four neighbors. Cells that are part of empires never attack other cells in their empire. A successful attack occurs if the power of the attacking cell is sufficiently greater than the power of the defending cell.

The *power* of a cell depends on 3 factors:

- Power increases with the asabiya throughout the empire, indicating the importance of social solidarity in military victory.
- Power increases with the total size of the empire, indicating the total resources available to the military.
- Power decreases with the region's distance from the empire's center, indicating the time and energy for communications and resources to travel.

Formally, the power of a cell *i* belonging to empire *j* is calculated as follows:

$$P_i = A_i \acute{S}_i \exp\left[\frac{-d_{ij}}{h}\right]$$

where d_{ij} is the distance from the cell to the imperial center (the geographical center of the empire) and *h* is a parameter determining how rapidly power declines with increasing distance from the imperial center. If the difference between the attacker and defender cells is greater than delta-p, the defending cell is taken over, and its imperial index

Finally, an empire can collapse completely if its average asabiya becomes too low. If $\acute{S}_i < S_{crit}$, the empire is dissolved and all its cells become independent chiefdoms.

Our primary question for this model is: do we observe cycles of the rise and fall of empires that look anything like those seen historically? This model involves considerably more coding than previous models. As with the

previous models, I have provided the code; this one may take a little more work to understand than the others. In general, some computational models are quite complicated (with some considerably more complicated than this one!). It's important to consider how many assumptions are necessary, and to always make the model as complex as necessary, but not more so. Due to the complexity of the model, I won't go through most of the parameters, but rather leave that an exercise for the curious student (I will adopt parameter values used by Turchin 2003 unless noted otherwise).

SETTING UP THE MODEL

- *allow-collapse?* switch. This controls whether empires collapse if their average asabiya gets too low, or instead diminish only through conquest.
- *power-decline* slider. This is *h* in the model. Power declines more rapidly with smaller *h.*

INITIALIZATION

- Each patch is initialized as a chiefdom with a baseline *asabiya*.
- The four patches in the corner of the grid are consolidated into an empire.

DYNAMICS

- Update each cell's asabiya based on its proximity to a frontier.
- Calculate stats of area, asabiya, and power for each empire.
- Calculate the power for each cell.
- ASK PATCHES
    - Consider attacks on all neighbors. Update empires accordingly.
- Collapse any empire in which asabiya has fallen too low.

PLOTTING:

- Empires will be assigned colors, which will spread on the spatial grid.
- Plot the area of each empire over time as it appears.

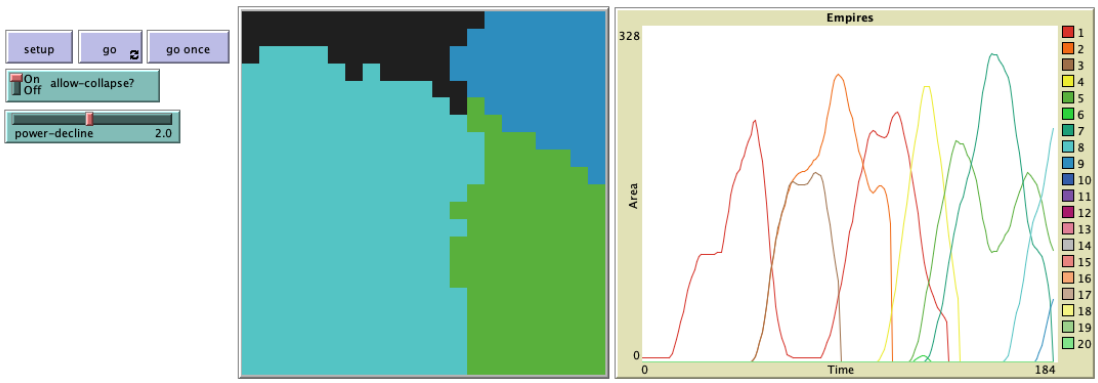

RESULTS

- When we run the model, the first thing that we observe is that, after initial growth by the first empire, it is eventually attacked by a small chiefdom and diminishes as other empires rise. We then get very clear cyclical behavior as empires rise and fall.
- Changing the value of *h,* the *power-decay*, has strong effects on political stability. When h is smaller, power decays very rapidly with distance from the imperial center, which should curb the growth of empires. It does this, however because it also does not allow empires to grow very powerful, it also increases cultural stability. We get the emergence of several moderately-size empires that are relatively stable with minor border skirmishes. (h = 0.7). This indicates that technologies that improve travel and communication over long distances can also lead to larger but less stable societies.

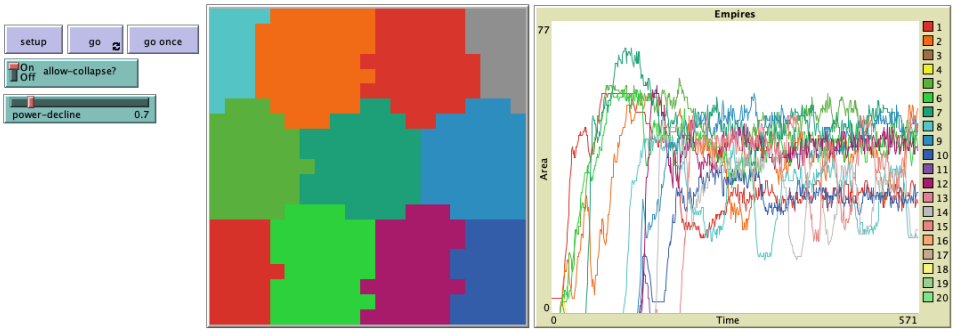

- When *h* gets larger, power declines more slowly and empires can grow quite large. As they grow to equilibrium, however, they become highly unstable, and we get very rapid change over, and sometimes population collapse. For example, when h = 3, we see the cyclical rise and subsequent collapse of very large empires, sometimes leading to the complete collapse of all existing empires (the model does not have a mechanism for the formation of new empires when none exist). This further indicates how mobility and communication technologies might contribute to sociopolitical stability.
- If we no longer allow empires to collapse because of low asabiya (back to h = 2), we still get cycles of growing and shrinking empires, but it is often the same few empires competing over and over, with very few new empires emerging. This indicates that, if the other assumptions of the model can be considered valid, that the collapse of large empires due to non-military factors is essential to understanding historical dynamics.

Does the model show that the Metaethnic Frontier theory is right?

Turchin (2003) notes a strong similarity between the model output and the actual historical dynamics of the empires of East and Central areas over a 600 year period. I've reproduced his Figure 4.4 below This is a nice demonstration of possibility: our assumptions generated the sort of dynamics seen empirically.

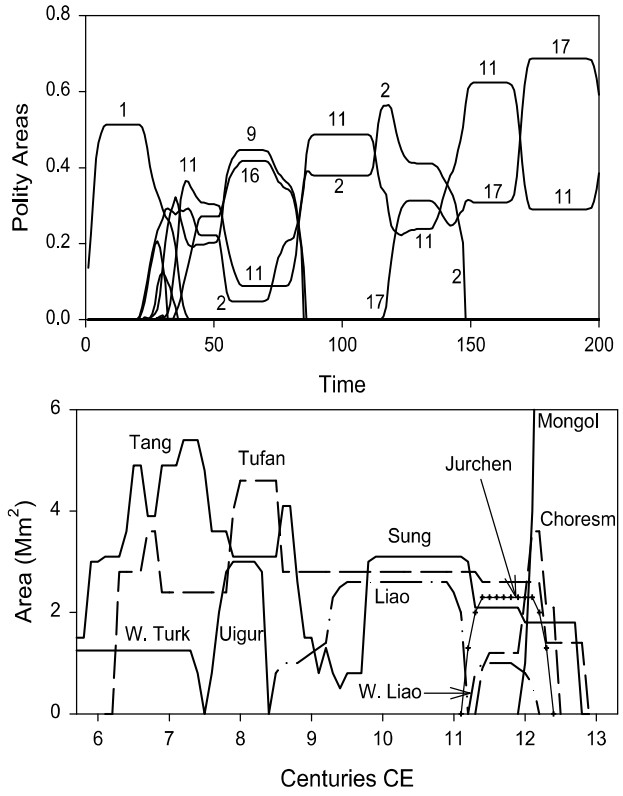

Figure 4.4 (a) Dynamics of the spatial asabiya-area model. Each curve depicts the territorial dynamics of simulated empires (polity area is expressed as a fraction of the total arena occupied). Model parameters are: $r_0 = 0.2$, $\delta = 0.1$, $h = 2$, $S_{\text{crit}} = 0.003$, and $\Delta_p = 0.1$. Numbers associated with trajectories are imperial indices of the polities. (b) Expansion-contraction curves of areas for polities in East and Central Asia, 600-1200 CE. (Data from Taagepera 1997: Appendix).

That said, proof of possibility is not proof of necessity. The computational social scientist Josh Epstein has said, in promoting the use of models like this one, that "If you didn't grow it, you didn't explain it." I agree that this is a strong argument in favor of modeling. However, the inverse does not necessarily follow. If you did grow it, you have not necessarily explained it, only failed to reject a possible explanation. Turchin himself is very clear on this point, and it's worth quoting him at length here:

> "The qualitative similarity between historical polity trajectories and the simulated ones does not, of course, constitute any 'proof' that the theoretical and empirical dynamics are driven by the same mechanisms. In general, for any potential pattern there are an infinite number of mechanisms that can produce it: the mapping between mechanisms and patterns is many to one... Nevertheless, the

observation that several features of the model's output match the observed dynamics is, at the very least, an encouragement to further theory developing and testing." (Turchin 2003, p. 71).

Obviously, this model is far from a complete theory of historical dynamics. However, the simple model shows us that the metaethnic frontier idea probably represents a plausible theory, and can indeed generate patterns that approximate what we see in history, at least in some places and at some scales. More generally, social dynamics are complicated, and models like this one give us a foothold to translate our verbal theories into concrete dynamical ideas.

Further directions

- The evolution of contagion. If one posits a simple mixed population of near-infinite size and considers the evolution of pathogen virulence, one must conclude that pathogens should evolve to be as contagious as possible – that is, having the highest possible transmissibility, as they will spread farther and faster. However, this conclusion does *not* hold in a spatially structured population of finite size. By allowing traits like transmissibility to mutate and evolve, Goodnight et al (2008) showed that the consideration of population structure in the spatial model yields the prediction that strains that are overly transmissible will soon exhaust their local supply of hosts and perish, whereas more moderately transmissible strains will persist.
    - Goodnight C, Rauch E, Sayama H, De Aguiar MAM, Baranger M, Bar-Yam Y (2008) Evolution in spatial predator-prey models and the 'prudent predator': The inadequacy of steady-state organism fitness and the concept of individual and group selection. *Complexity* 13(5): 23–44.
- Getting more complicated with historical cycles. Turchin and colleagues have continued to develop theories of historical dynamics, adding factors including environmental factors like oceans and mountains, technological innovation, and conflict among elites to refine their theory and better fit it to available historical data. The models have also inspired more focused collection of historical and archaeological data with an aim to better parameterize dynamic models.
    - Turchin P, Currie TE, Turner EAL, Gavrilets S (2013) War, space, and the evolution of Old World complex societies. *Proceedings of the National Academy of Sciences* 110: 16384–16389.
    - Turchin P (2016) *Ages of discord: A structural demographic analysis of American history*. Beresta Books.
- Growing artificial societies. There have been a number of efforts to model complex historical dynamics with complex, realistic agent-based

models. One of the first was Epstein and Axtell's 1996 book *Growing artificial societies*, which includes a careful study of their stylized Sugarscape model. This has been followed by many studies by social scientists and archaeologists to use models to better understand our past.

  - Epstein JM, Axtell R (1996) *Growing artificial societies*. MIT Press.
  - Epstein JM (2006) *Generative social science: Studies in agent-based computational modeling*. Princeton University Press.
  - Crabtree SA, et al. (2017) How to make a polity (in the central Mesa Verde region). *American Antiquity* 82: 71–95.

- Autocatalysis. The models we explored in this unit share principles with a classic sort of model in the physical and biological sciences: those of autocatalytic processes. Many of these models consider substances that are characterized by local activation and long-range inhibition. That is, the substance is autocatalytic in close range but faces (indirect) inhibition by increased distance from the central zone of influence. Such forces can generate churning cycles and pattern formation. Such models are widely used to explain processes of chemical and biological pattern formation.
  - Kondo S, Mura T (2010) Reaction-diffusion model as a framework for understanding biological pattern formation. *Science* 329: 1616–1620.
  - Jones J (2010) Characteristics of pattern formation and evolution in approximations of *Physarum* transport networks. *Artificial Life* 16: 127–153.

Exercises

- *Space is the place.* Play with population density in the host pathogen model. What if a proportion of cells are initially empty. Does this affect how often the pathogen dies out? What is some cells are not only empty, but unfillable, representing barriers in physical or network space?
- *Evolving contagion.* Adapt the host-pathogen model so that transmissibility is an evolvable trait (see Goodnight et al. 2008 for suggestions how to accomplish this). Do pathogens evolve to be as transmissible as possible? Why or why not?
- *Making more room.* In the metaethnic frontier model, we used a 21 x 21 grid. How does the overall size of the territory (in terms of habitable regions) influence the spread of empires? What if make the grid considerably large, like 51 x 51? How does this effect our conclusions about the number of empires at one time, and the persistence of those empires?

- *Thinking through assumptions.* Consider the metaethnic frontier model. Describe the model's assumptions. Why do those assumptions lead to cycles (or fail to do so)?
- *Possibilities.* Imagine you had advanced the metaethnic frontier model so as to be able to accurately model the rise and fall of modern states. How would you calibrate that model? Would it be possible to make concrete predictions?