UC Riverside
Cliodynamics

Title
On the Life Cycle of Empires

Permalink
https://escholarship.org/uc/item/8p0176zr

Journal
Cliodynamics, 16(2)

Author
Ciliberti, Stefano

Publication Date
2025-12-25

DOI
10.21237/C7clio.53812

Data Availability
The data associated with this publication are available at:
https://empires.wnvermeulen.com/

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, available at
https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

On the Life Cycle of Empires

Stefano Ciliberti
University of Pisa (Italy), and EPHE (Paris, France)

We provide empirical evidence that the distribution of empires' lifetime is exponential with a typical time-scale of 300 years. After introducing and computing a proxy measure for the dimension of a large polity we interpret the data by means of a simple dynamical model for the empires' evolution. The resulting theory suggests that the behavior of a long-lived empire is qualitatively different from a short-lived one. They all experience an initial expansion phase but while the former goes through a contraction period over the last 20% of its life, the latter stays essentially unchanged after reaching its maximum size. In both cases, a sudden collapse occurs when the empire's size is still close to its own historical average.

Introduction

In his now outdated monumental work (Toynbee 1962) Toynbee attempted to analyze the rise and fall of civilizations through a systematic and comparative approach. He argued that the development of civilizations is driven not by race, environment, or economics, but by how societies respond to challenges – both internal and external. Often criticized for his lack of empirical rigor and the philosophical overtone, not to mention the ideologically eurocentric view of the very concept of civilization, Toynbee's work had the merit to show that a large-breadth ambitious study could inspire macro-historical thinking like world history, comparative history, or even international relations. His methodology has since been superseded by more complex, networking or systemic models. Examples include Wallerstein's world-systems theory (Wallerstein 1974), where it is proposed that global economic systems should be the real object of study of historians, instead of individual civilizations or nations; Diamond's environmental history (Diamond 1997), based on the claim that civilizations rise or fall largely due to environmental constraints, geography, and access to resources; Braudel's (Braudel 1949) focus on "structures", like geography, trade, and material life, not just elites or politics, along with the cogent idea that history unfolds on multiple timescales (events, conjunctures, and *longue durée*); or Turchin's cliodynamics (Turchin 2016), where the core idea is to use data modeling and historical patterns to study cycles of stability and crisis in societies.

Corresponding author's e-mail: stefano.ciliberti@gmail.com

Citation: Ciliberti, Stefano. 2025. On the Life Cycle of Empires. *Cliodynamics* 16: 1–12.

The utility and coherence of the concept of empire in comparative historical studies has been questioned, among others, by Reynolds about twenty years ago (Reynolds 2006), who called for a more cautious and nuanced use of the term "empire" in historical writing. Rather than assuming similarities across empires, scholars should focus on specific institutions, relationships, and power structures, and avoid using "empire" as a catch- all. The risk of using an imprecise and often anachronistic label to group a wide range of political formations (e.g. the Roman, British, and Mongol Empires, sharing only few meaningful characteristics) may lead to establish false analogies. Many so- called empires differ radically in structure, ideology, economy, and territorial control: in pre-modern contexts especially, sovereignty was often overlapping, negotiated, and unclear. Unlike the centralized, absolute control that the traditional idea of empire often implies, these large political entities did not always dominate their peripheries: they frequently relied on alliances, client states, and indirect rule.

In a recent two-volume book Bang et al. 2021 that brings together several scholars to examine empires across history through a comparative, thematic, and interdisciplinary lens, the notion of a singular fixed definition of empire is rejected and pragmatically replaced by a heuristic category: a flexible tool for comparison, rather than a rigid classification. Empires are recognized as historically variable, appearing in different regions, periods, and forms, yet always defined by scale, hierarchy, diversity, and asymmetry of power. Comparative studies are welcome, as long as the contextualization is intended as a counterbalance to the itch to force everything into a single model.

In this paper our focus is not on economy, ideology, or political structures, but rather on time and space. Following Reynolds, and in line with the operational approach in (Bang et al. 2021), we lump together under the term "empires" the political entities able to enforce their laws and consisting of one ruling polity (the metropolis) and one or more subordinate polities (peripheries or colonies): at odds with federal states, for instance, where the different components are supposed to be equal. This working definition, while sweeping under the carpet many key details on the nature of sovereignty and territorial control, yet allows to talk about when and where such a political hegemony is established. In particular, we are interested in the lifespan of empires and whether or not one could clearly identify "phases" based on their territorial extent.

The problem of understanding the life cycle of empires has attracted the attention of different scientific communities over the last few decades. Some scholars have focused on the expansion phase, the route to large territorial domains: Turchin (Turchin 2009), for instance, proposed a model where the interplay between two populations, nomadic and settlers, generates a self-sustaining process leading to the creation of a strong political entity supported by

a structured military power. In the majority of the studies, however, the goal is to try and understand how, when and why empires collapse: from Gibbon (1932) to Tainter (1988) and Diamond (2005) a large effort has been made on the contraction, fragmentation, and disintegration phases; usually a mix of both exogenous and endogenous elements are held responsible for the collapse of complex entities. In most of these approaches, providing one or more examples where the theory works well seems to be the most effective way to get the message across – rather than backing it with data spanning over different time periods and geographical areas.

In a recent paper Arbesman (2011) has focused on the distribution of empires' duration by taking advantage of the data compiled by Taagepera for his seminal studies (1978; 1979). This dataset contains approximate size and dates for 41 empires spanning the period between 3000 BCE to 600 CE. The author finds that the distribution of lifetimes is, to a good approximation, exponential.

The first goal of this paper is to address the same question by using a much larger database to check the validity and robustness of Arbesman's results. We then propose and compute a stand-in measure of an empire's dimension and follow its evolution from birth to death. A simple mathematical model is used to fit the data and describe the typical behaviour of an empire throughout its whole life cycle. Our findings can be thought of as evidence for a simplified yet universal pattern of empires' life cycle and collapse, where a distinction between short-lived and long-lived polities appears to be relevant.

In the next section we describe the data and present the results of our analysis about lifetime and size; the dynamical model is also introduced and used to interpret the empirical findings. A summary discussion is presented in the final section, along with some suggestions for future works.

2 Data, model, and results

In order to support the idea that the very existence of empires facilitated trades within their territories, Gokmen et al. (2020) collected data on rise and fall of global empires over the last five thousand years. The database contains 168 different political entities, the approximate start/end years, and some more details including the number of (modern) countries included in each empire at a given time[1]. The sources for these informations are diverse and range from specialized books to historical atlases or online encyclopedias like Britannica and World History. The very definition of chronological boundaries depends on some considerations which may come across as controversial: for instance, the Roman and Byzantine Empires are both treated as one long political entity, while every

[1] https://empires.wnvermeulen.com; www.britannica.com; www.worldhistory.org

Chinese dynasty is different. We believe that this should not be a major problem for our statistical study because the results, as we will see below, turn out to be quite robust.

The oldest empires recorded in the database go back to the third millennium BCE (Elam and the Akkadian Empire). The five longest are the Pandyan dynasty in south India, Elam in Iran (both around 2000 years), the Roman and Byzantine Empires, and the Republic of Venice, all three just above one millennium. The shortest-lived political entities are the First Mexican Empire, queen Zenobia's Palmyrene kingdom, the Third Reich and the Korean Empire, all of them lasting only a few years.

2.1. Duration

The empirical average duration is 298 years, and if we remove the outliers (i.e. the durations larger than 1500 years) that average goes down to 277. Without resorting to sophisticated statistical firepower, it is reasonable to assume that the unbiased estimate of the mean value lies between these two numbers, which are very close anyway. The standard deviation computed from the data is 310 years, roughly equal to the average: a first clue that the distribution might be exponential. It is also worth noting that the median is 210 years, which means that at least half of the empires lasted more than two centuries; moreover, the median for an exponential distribution is $\ln(2)\times$ mean, which gives $\simeq$ 207 years in our case.

Synthetic indicators such as mean, standard deviation or median might be extremely useful to get a first feeling of the data but a deeper insight requires a look at the full distribution. In Fig. 1 we show the complementary cumulative distribution of the duration of the empires, that is our estimate of the chance $\mathrm{P}\,(T \geq t)$ of observing a duration T larger than (or equal to) a given time t. Looking at this quantity instead of the simple empirical frequency is common in data analysis because it ensures a significant noise reduction in the visualization. On the x-axis the duration t is measured in years, and the probability on the y-axis (on a logarithmic scale) is computed by counting for each t the number of entities which lasted as much as or more than t: this number is then divided by the total number of observations in such a way that as the duration increases the probability goes from one to zero (or, in our finite size real-case, to 1/168).

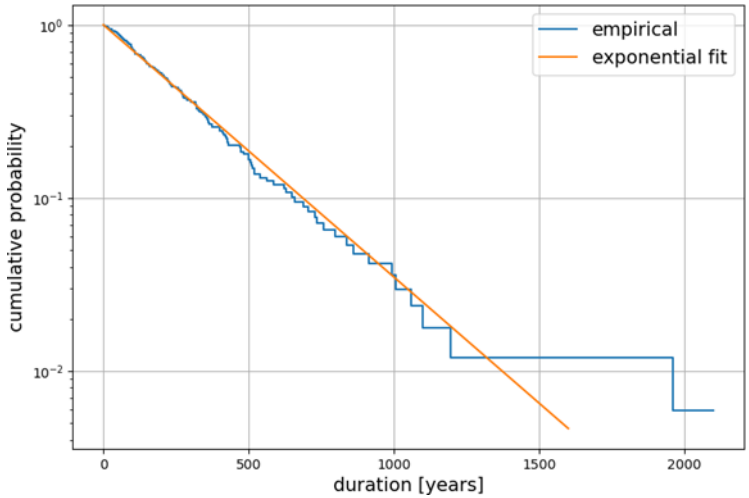


Figure 1. The cumulative distribution of empires' lifetime, along with the exponential fit. Note the log-scale on the y-axis.

The straight line on the same panel is the exponential fit $\mathrm{P}(T \geq t) = e^{-t/\tau}$, where the time-scale $\tau = 298$ is equal to the empirical average. The fit works well on the data over two orders of magnitude, a remarkable feature. Pushing further the analysis we have performed a Kolmogorov-Smirnov test (which is known to work well even with relatively small sample sizes) to compare our data sample with the theoretical exponential probability. We obtain a *p*-value of 0.71, well above the standard threshold of 0.05, a strong evidence in support of the exponential model[1].

Given these findings, we can safely conclude that the data are compatible with an exponential model for the duration of empires, thus confirming Arbesman's results. Nevertheless, he finds $\tau \simeq 220$, which is somehow smaller than our value.

As already pointed out by Arbesman himself and others, the exponential distribution has the memoryless property, which says that the likelihood of the next event does not depend on any past information. In our case, it means that the probability of collapse of an empire does not depend on its age: it can happen at every time, and having survived for long does not make a political entity more stable. This is not very intuitive as one would expect a centuries-old empire to be more resilient to, say, exogenous perturbations or internal crisis: but this argument is flawed because an empire is far from being an isolated system. Instead, it is usually embedded in a complex environment and competing with "conspecifics" for the same resources. This remark, along with the analogy with

[1] To be precise: a high *p*-value does not prove that that data is exponentially distributed, but it points at the lack of statistically significant evidence against that assumption.

biological species, led scholars to stress the role of what has been called the "Red Queen" effect (Van Valen 1973), where civilizations constantly adapt to a changing and strongly competitive environment.

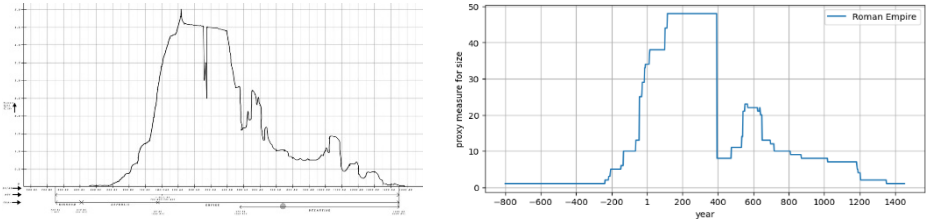


Figure 2. The example of Roman Empire: A comparison of our proxy measure of size (the number of countries included in an empire, right panel) with Taagepera's chart (Taagepera 1979) (left). The similarity is strong and suggests that we can move forward with our experiment. Note that here we have aggregated Roman and Byzantine Empires, the purpose being the comparison with Taagepera's plot.

2.2 Size

In order to gain a deeper insight into the dynamics of empires' cycles we have built a proxy measure for their size by simply counting the number of modern countries they encompass at a given time. This can be grossly inaccurate in some cases, but a priori it does not contain a systematic bias; most importantly, we are interested here in the average behaviour and the large number of entities studied (168) is probably sufficient to dilute away idiosyncratic effects: we will come back to this point below.

A first successful check (see Fig. 2) is done by simply comparing our proxy metrics with some previously published results (Taagepera 1979; 1997) where an attempt was made to report the areas in square kilometers for several empires over their whole existence. The comparison works relatively well in most cases (like the Ottoman, Mongol, or Holy Roman Empires) while is less convincing in others. As no systematic under- or over- estimation has been observed by going through this exercise, we expect the individual errors to offset when computing averages. It may also be tempting to go beyond the qualitative level and infer, on the basis of a few particularly well-matched cases, that 10 countries roughly correspond to (say) 1 million square kilometers, but although this mapping seems reasonable, we do not deem it statistically significant.

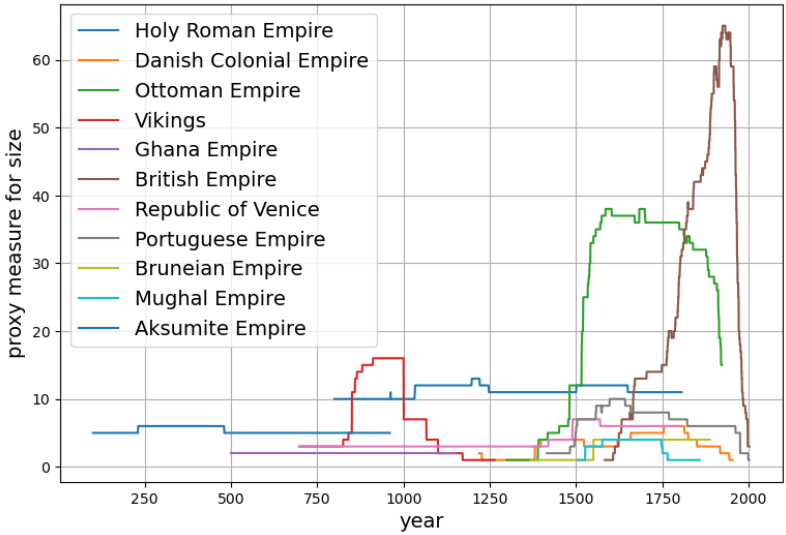


Figure 3. A few examples of the size proxy measure.

Other examples of our metrics are plotted in Fig. 3 and show that the life of an empire can be very diverse. It can start from being very small, grow big and then shrink slowly and gradually (like the Roman Empire), or it can be already big when created and keep a comparable size for the rest of its life (e.g. the Holy Roman Empire); actually, all the situations in between these two extremes are also possible (see for instance the Republic of Venice).

Our question now is: what is the average behaviour of size vs time during an empire cycle? And how typical is this average? These points are addressed in Fig. 4. As we want to aggregate data for all the empires, we define for each of them a "normalized age", that is a time variable going from 0 (start date) to 1 (end date). Once we do this operation, which comprises resetting and rescaling, we can safely combine different empires and compute the average and standard deviation of the size over the lifespan: the result is shown in the left panel of Fig. 4. A few comments are in order:

- The average size is not close to zero at the beginning, and it is roughly 2/3 of the maximum size (recall that the unit on the y-axis is somewhat arbitrary so we cannot easily translate it into square kilometers)

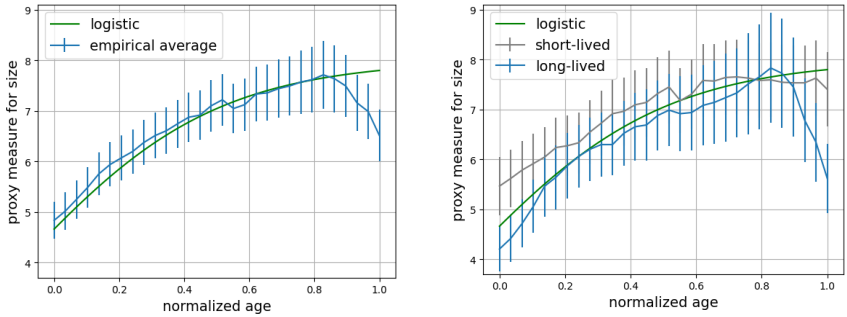


Figure 4. Left: the average size of empires is computed as a function of their normalized age; the logistic fit is also shown. Right: the average size is computed separately for short-lived and long-lived empires. See text for more explanations.

- The average curve shows an increase for 80% of its lifetime (expansion), after which the size decreases (contraction).
- On average, the size of the empire is not small before it collapses. As a matter of fact, the last point in the chart is very close to the average of all the previous numbers.
- The errorbars, here computed in the textbook way as the empirical standard deviation divided by the square root of the number of data points, are relatively small. This fact implies that the average behaviour is representative of the typical one.
- If we were to go back to natural time units, we should multiply the x-axis by the average duration, which we have seen before to be about 300 years; this means for instance that the peak happens, on average, after two and a half centuries from the birth.
- The data are well fitted by a logistic function up to the maximum, but this fit fails at describing the last 20% of the curve: we address this point below.

This result would lead us to interpret the average life cycle of an empire as one entity which is already big when created, it grows for some time, then starts decaying to some extent and collapses when its size is still close to its own historical average.

How robust is this interpretation? In order to test the reliability of this idea we have ran two numerical experiments: in one case we have split our data set into different periods (i.e. before or after 500 CE), in the other scenario the split is based on the duration (i.e. smaller or bigger than the average). While the ancient vs modern split does not lead to two significantly different curves, the short-lived vs

long-lived separation shows something interesting (right panel of Fig. 4). The pattern of empires lasting less than three centuries is monotonous: they grow until they reach their maximum (typically 50% higher than their dimension at the start), then remain at that level until their death – which we can suppose sudden and unexpected because of the lack of visible signs of contraction. On the other hand, long-lasting empires have a similar pattern in the first part but decay in a significant way before their death: this piece of the curve in the global average (left panel) thus mainly comes from the empires that last more than the average duration of three centuries.

2.3. Modeling

The last point to discuss about this chart is the logistic fit, which calls for an explanation. We can build a model for the dynamics of an empire's size (let's call it $X^{(k)}(t)$, k being just the label of the empire) by assuming that its yearly variation is driven by three factors: *i*) a natural expansion process happening along the current boundaries (thus proportional to the actual size); *ii*) the pressure from the surrounding states which counterbalances the previous factor and is proportional to the size of both the states sharing a boundary; *iii*) exogenous factors completely independent from the empire size. In formulas:

$$\frac{dX^{(0)}}{dt} = \alpha X^{(0)} \left(1 - \gamma \sum_{k\in 0} X^{(k)}\right) + \eta, \qquad (1)$$

where α is the natural rate of growth of an empire in absence of its neighbours: a ballpark figure for this parameter should be one percent because a too large yearly variation of the size is unlikely to be compatible with some stability[1]; γ represents the average weight of the neighbours, whose impact is proportional to their own size (the sum runs over all the nearest neighbours k of the empire 0); η is a stochastic external noise whose time scale we expect to be large and which can be taught of as the exogenous random factor.

Given that there is nothing special about empire (0) with respect to all the others in the above equation, we can assume that the dynamics for the average of all of them is well approximated by the following:

$$\frac{d\langle X\rangle}{dt} \approx \alpha\langle X\rangle(1 - \gamma\langle X\rangle) + \langle\eta\rangle, \qquad (2)$$

within a mean-field kind of approximation on the right-hand side: $\langle X^{(i)}X^{(j)}\rangle \simeq$

[1] An empire can definitely double in size in one year ($\alpha \simeq 1$), but if this happens every year the process is probably too fast to be stable.

$\langle X^{(i)} \rangle \langle X^{(j)} \rangle$. If we now suppose that the average impact of external factors $\langle \eta \rangle$ can be neglected then we are left with a classical differential equation whose solution is given by the logistic function:

$$\langle X(t) \rangle = \frac{\gamma^{-1}}{1 + e^{-\alpha(t-t_0)}}, \tag{3}$$

where γ^{-1} is the $t \to \infty$ limit of $\langle X(t) \rangle$, α^{-1} is the time scale of the process, and t_0 is the first observation time. This function can now be used to fit the data in Fig. 4 and the parameters used are $\gamma^{-1} \simeq 8$, $t_0 \simeq -0.1$, and $\alpha^{-1} \simeq 0.3$: these last two numbers can be easily translated into natural units after multiplying them by the average lifetime measured in the previous section, yielding to $t_0 \simeq -30$ years[1], and $\alpha^{-1} \simeq 90$ years. This step clears the way for an interpretation: such a negative value for t_0 hints at a narrative where some growth dynamics has already occurred over the previous generation when we start observing the process; as for the second parameter, $\alpha \simeq 1/90 \simeq 1.1\%$ per year, it confirms our intuition about the order of magnitude of the natural expansion of an isolated empire, and indicates that the most significant increase in size (the acceleration in the expansion phase) happens over a period slightly smaller than a century.

It is important to stress that the logistic function (eq. 3) is a solution to the model only when the average impact of external factors is neglected ($\langle \eta \rangle = 0$ in eq. (2)). This is not a realistic assumption but as the fit seems to work well on short-lived empires (see the right panel of Fig. 4) we can conclude that the time-scale of this exogenous factor is longer than a few centuries: in other words, the collapse of a large empire is more likely to be induced by a massive invasion, a major climate crisis or a significant natural catastrophe – in any case a relatively rare event.

3 Discussion and perspectives

In this paper, we address fundamental questions regarding the temporal and spatial dynamics of empires: How long do empires last? How large do they grow relative to their initial and final states? Can we identify distinct phases of expansion and contraction within their life cycle? And, from these patterns, can we infer the most probable causes of collapse?

We do not lay claim to settling such complex problems, but our statistical analysis highlights a "typical" pattern of behavior. Based on the 168 political entities studied, we confirm that the lifespan of empires follows an exponential distribution: older empires are not necessarily more stable than more recent ones.

[1] An empire can definitely double in size in one year ($\alpha \simeq 1$), but if this happens every year the process is probably too fast to be stable.

Using territorial extent as an indicator allows us to reconstruct a typical life cycle: an initial phase of expansion – sometimes preceding the formal constitution of the political entity – is followed by a period of stability at maximum size, which may end abruptly. When this phase lasts for several centuries without external disruption, a slight decline is generally observed before disappearance, often triggered by an exogenous event. On average, empires endure for about three centuries, reach their peak after roughly one century, and then experience either rapid collapse or keep alive longer, shrink to a mild extent, and die while still big.

Mathematical models capturing this dynamic remain rare and simplified, often isolating a single factor (Aldous 2009). Developing a model that combines territory, population, and wealth would provide a deeper understanding of the mechanisms underlying imperial growth and collapse. Our model represents a first, imperfect step in this direction: it would be relevant to incorporate interactions between empires, processes of fusion and fragmentation, and to connect these dynamics with Turchin's theories, in which the interaction of distinct populations can give rise to new empires.

In summary, despite the complexity and variability of trajectories, statistical analysis allows us to identify an average imperial life cycle with measurable timescales: initial growth, peak, and then decline or collapse. This perspective provides a reference framework for comparative study and paves the way for more refined models that integrate both internal and external factors, thereby contributing to a better understanding of the historical dynamics of empires.

Acknowledgment

I am deeply grateful to Stéphane Pequignot, Alma Poloni and Andrea Puglisi for their support and comments to this work

References

Aldous, David J., J. R. Ong, and W. Zhou. 2009. Empires and percolation: stochastic merging of adjacent regions. *Journal of Physics A: Mathematical and Theoretical* 2: 025001.

Arbesman, Samuel. 2011. The Life-Spans of Empires. *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 44 (3): 127-129.

Bang, Peter F., Christopher A. Bayly, and Walter Scheidel. 2021. *The Oxford World History of Empire, Oxford*. Oxford University Press.

Braudel, Fernand. 1949. *The Mediterranean and the Mediterranean World in the Age of Philip II.* Armand Colin editeur.

Diamond, Jared. 1997. *Guns, Germs, and Steel: The Fates of Human Societies.* W.W. Norton.

Diamond, Jared. 2005. *Collapse: How societies choose to fail or succeed*. New York: Penguin.

Gibbon, Edvard. 1932. *The decline and fall of the Roman Empire*. New York: Modern Library.

Gokmen, Gunes, Wessel N. Vermeulen, and Pierre-Louis Vézina. 2020. The imperial roots of global trade. *Journal of Economic Growth* 25: 87–145. https://empires.wnvermeulen.com/

Reynolds, Susan. Empires: a problem of comparative history. *Historical Research* 79 (204): 151–165.

Taagepera, Rein. 1978. Size and Duration of Empires: Growth-Decline Curves, 3000 to 600 B.C. *Social Science Research* 7 (2):180-196.

Taagepera, Rein. 1979. Size and Duration of Empires: Growth-Decline Curves, 600 B.C. to 600 A.D. *Social Science History* 3 (3-4):115-138.

Taagepera, Rein. 1997. Expansion and Contraction Patterns of Large Polities: Context for Russia. *International Studies Quarterly* 41(3): 475-504.

Tainter, Joseph A. 1988. *The collapse of complex societies*. Cambridge: Cambridge University Press.

Toynbee, Arnold J. 1962. *A Study of History,* Oxford University Press.

Turchin, Peter. 2009. A theory of formation of large empires. *Journal of Global History* 4: 191–217.

Turchin, Peter. 2016. *Ages of Discord: A Structural-Demographic Analysis of American History*. Beresta Books.

Van Valen, Leigh. 1973. A new evolutionary law. *Evolutionary Theory* 1: 1–30.

Wallerstein, Immanuel. 1974. *The Modern World-System. Academic Press.*