Cliodynamics: The Journal of Quantitative History and Cultural Evolution

Repeated Demographic-Structural Crises Propel the Spread of Large-scale Agrarian States throughout the Old World

James Bennett
University of Washington

Abstract

I investigate the geographical consequences of demographic-structural dynamics using a spatially resolved agent-based model of agrarian empires in several Old World regions between 1500 BCE and 1500 CE. I estimate and bound key model parameters from two historical datasets. Although several very large-scale polities (e.g., Roman, Persian, Tang empires) do not arise and certain geographical expansions occur at different times, overall the model suggests that factional civil wars, the result of repeated internal demographic-structural crises, can substantially account for the spread of large-scale agriculture throughout the Old World after the Bronze Age.

Introduction

Demographic-structural theory (Goldstone 1991) forms the core of much recent investigation of large-scale historical events, providing a framework for understanding the corroding impact of demographic pressures on cohesive political and social structures, accelerating wealth inequality, economic stagnation, intra-elite competition, and finally triggering rebellion and civil war. Turchin and Nefedov (2009), in particular, provide a succinct review of the theory and analyze several historical periods in order to evaluate the theory's utility, concluding that it offers a compelling account of the growth and decay of states ('secular cycles') observed in the historical record.

Mathematical formulations and models of this theory (Turchin 2003, Turchin 2009, Nefedov 2013) demonstrate that small sets of coupled non-linear differential equations capture the major qualitative behaviors of the theory and provide quantitative predictions of wealth and population growth, the duration of stagnation phases, and the dynamics of state collapse. However, most formulations have been applied only to abstract, individual, and isolated polities with an implicitly fixed carrying capacity (an environment's maximum

Corresponding author's e-mail: jsb11@u.washington.edu

Citation: Bennett, James. 2016. Repeated Demographic-Structural Crises Propel the Spread of Large-scale Agrarian States throughout the Old World. *Cliodynamics* 7: 1–36.

sustainable population) or spatial size with some parameterizations inspired by historical events and trends.

Agent-based model investigations of polity formation and collapse, on the other hand, have employed different spatial models, not all inspired by demographic-structural ideas, in both realistic and abstract geographies. Artzrouni and Komlos (1996) proposed a model (here 'Artzrouni') of state formation and applied it to the European theater between 500 CE and 1800 CE. Their model focused on the projection of military force required to annex and defend territory. They observed that polities with higher marchland advantage—those enjoying lowered defensive costs afforded by mountains or sea-based physical boundaries—faired particularly well. This advantage led to larger polities that stabilized quickly in the western half of Europe (thanks to the easily defended Mediterranean, the English Channel, and the North Sea) and left many smaller states to contend for a longer period of time in the plains of eastern Europe. Their model, however, did not address any internal collapse dynamics within polities, nor did it attempt to predict any demographic consequences.

Turchin (2003, Appendix 2, here 'A2') investigated an 'asabiya'-driven spatial model of polity rise, expansion, and fall. Motivated by meta-ethnic conflict analysis, the model posited that a polity's ability to annex and defend territory depended on the sum of each region's level of cooperation or group cohesion (*asabiya* (Ibn Khaldun 1958)) toward their polity and against their neighbors. Asabiya increased logistically for frontier regions that bordered other polities, otherwise decaying exponentially for interior regions distant from border threats. According to this model, large polities collapsed and splintered into individual regions if the polity's combined asabiya dropped below a fixed threshold.

Similar to Artzrouni and Komlos, Turchin's A2 model assumed a polity's force was based on combined regional asabiya and was projected to its political borders, falling off with distance. Although only exercised in an abstract, rectangular landscape, for certain parameter values the model demonstrated large-scale polity formation by single-region ('hinterland') polities on the borders of threatening extant empires and suggested that these new polities developed strength by first annexing weaker hinterland territories away from the original threatening empire ('reflux') before developing enough power to annex that eventually-weakened original empire's territories. While patterned after aspects of demographic-structural theory, the model neither attempted to predict population dynamics, nor employed it to influence empire behavior.

Turchin et al. (2013, here 'TCTG13') investigated the rise of large-scale agriculturally-based (agrarian) polities in the presence of enhanced military pressure from horse-riding steppe nomads. Modeling the acquisition, maintenance, and loss of abstract 'ultra-social' traits, increased cooperation within polities is reinforced and spread through surrounding areas in response to diffusing military prowess. Like the A2 model, the TCTG13 model posits that

polities arise and collapse as these (asabiya-like) ultra-social traits are acquired and lost. The TCTG13 model was able to account for up to 65% of the spatial variance observed in the location and timing of Afro-Eurasian historical polity formation between 1500 BCE and 1500 CE. However, there was no attempt to model any underlying demographic trends or their effect on polity formation.

Taken together, these previous modeling efforts neither employ demographic pressure to stress structural aspects of states in realistic geographies, nor investigate how these stresses might lead to state creation, expansion, and collapse accompanied by civil and inter-state war. Here, to address these issues, I develop an agent-based model to assess experimentally whether a more direct implementation of demographic-structural theory could account for the historical development of large-scale agrarian polities in the Old World between 1500 BCE and 1500 CE. In the model, each polity is an agent whose behavior is governed by internal demographic-structural forces as well as interaction with its environment and other polities. Starting from the small set of historical states present in 1500 BCE, simulations unfold independently of the actual historical record for 3000 years and are then compared to it.

In the following sections, I first review the geographical and historical data as well as the demographic model that I employed to develop and test my demographic-structural model. Next, I describe my agrarian demographic-structural model ('ADSM') in detail, then discuss its behavior in different regions of the Old World, evaluating its ability to predict both polity creation and collapse as well as their location, timing, and aggregate population sizes. I also discuss the model behavior's sensitivity to various parametric and boundary condition changes. The paper concludes with observations about the model, its shortcomings, and next steps.

The Historical Data and Demographic Model

Historical Data

As in Bennett (2015), I developed and evaluated my model employing two data sources. The first dataset, referred to here as the TCTG13 dataset, was developed by Turchin et al. (2013) to evaluate their model of the dynamics of ultra-social traits in the Old World. One part of the dataset describes the Afro-Eurasian continent as a gridded landscape of regions, encoding mean elevation and predominant 'biome type': desert, steppe, or agricultural. More fully described in their Supporting Information, regions were assumed to occupy 10 thousand km^2 (100 x 100 km). Each TCTG13 land region is connected to its immediate cardinal neighbors if they exist; regions are marked as littoral if any of their neighbors are an ocean or a sea.

The original TCTG13 dataset was supplemented by latitude and longitude coordinates for each region, rather than the ordinal values assigned for their study. This allowed the computation of the actual area of each region, which varies by (the cosine of) latitude with a border size of 111.32 km per degree at the equator. As it happens, the TCTG13 regions are, in fact, rectangular, stretching about 2 degrees of longitude and 1 degree of latitude and occupying about 25 thousand km^2 at the equator. In addition, several minor biome type and littoral encoding errors were also identified and corrected, none of which, like the area difference, would have any impact on their reported results, but are important to the results below.

The TCTG13 dataset also encodes the regional locations of all large-scale empires ('centralized macro-states', as they define it) present at each century between 1500 BCE and 1500 CE. Each encoded empire occupied at least 10 TCTG13 regions (250 thousand km^2, roughly the size of the contemporary Czech Republic or the historic Duchy of Burgundy). To compare my model results with their performance, I measured polity sizes using the same region-count metric even though the actual territorial size of the empires varied.

The second dataset encodes spatially and temporally resolved population estimates of Kaplan and Krumhardt (2008, here 'KK10'). This dataset contains yearly population estimates between 6050 BCE and 1850 CE. Based, in turn, on several well-known sources of historical population data (e.g., McEvedy and Jones, 1978), the population dataset is organized by modern, worldwide 'territories' (e.g., Paraguay, Minnesota, etc.). A separate territorial definition dataset, also supplied by Kaplan, provides latitude and longitude for each modern territory resolved to .08 degrees. Combined, the KK10 datasets permitted per-century estimates of population for selected rectangular subset landscapes of the TCTG13 Old World data, which I employed to evaluate and improve the model. Where the KK10 territories overlapped a TCTG13 subset region by less than 10%, their population estimates were discarded. TCTG13 subsets were chosen to avoid large population inconsistencies.

I focused on three subset landscapes in addition to the entire Old World: The European plus Middle Eastern theater ('Europe+ME'), the Indian subcontinent plus Middle Eastern region ('India+ME'), and east Asia including China, Japan, Korea, and South-east Asia ('Asia'). These subset region locations are indicated in Figure 1, which also shows the distribution of all agricultural (green) regions. Each of these study regions was chosen because they contain one or more initial empires required by my model (Egypt and the Hittite in 1500 BCE for Europe+ME and India+ME, the Shang for Asia, and possibly the Indus Valley civilization for India+ME). These empires and their descendants developed relatively independently of those in the other subsets, at least for the first millennium, and together these landscapes account for over 90% of the total population and empires of the Old World. Importantly, however, their spatial and population

growth dynamics are very different, apparently because of their very different geographies. (The northern steppe and African regions contribute very little to the agrarian empire and world population record in this time period.)

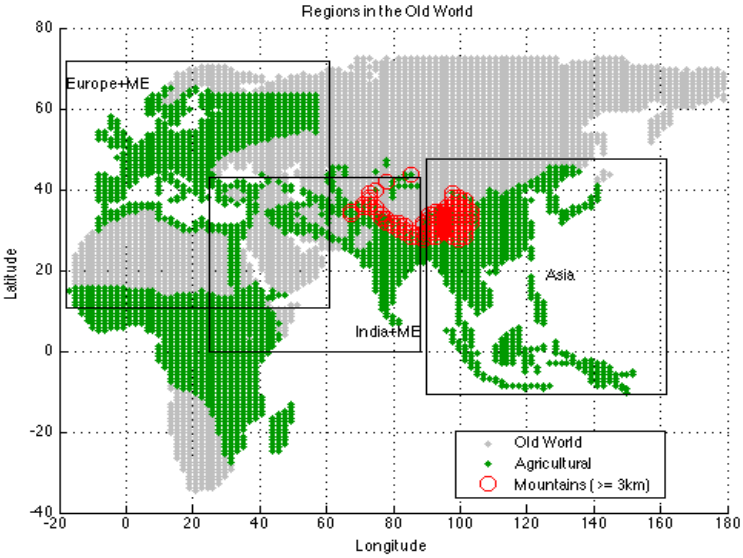

Figure 1. Agricultural regions in the Old World and the study subset landscapes.

There are roughly two major agrarian expansion periods (600 BCE to 100 CE and 700 CE to 1300 CE) between flanking periods of relatively little expansion. The waves of large-scale steppe nomadic empires that pestered the agrarian empires after 300 BCE are also clearly visible. These include the Xiongnu in 300 BCE, the Khazan in 700 CE, and the Golden Horde in 1300 CE. Two animations in Supplemental Materials, SM_Historical and SM_Historical_agricultural, show the spatial growth dynamics for all and agrarian-only historical empires, respectively.

Figure 3 shows the corresponding total population of the Old World to the nearest century from the KK10 dataset. The data show a weakly quadratic rise from 70 million to 390 million in this period with some major episodes of plague and famine clearly present in the record, notably the Justinian pandemic and Chinese famines and plagues beginning around 500 CE and the Black Death in the mid-fourteenth century.

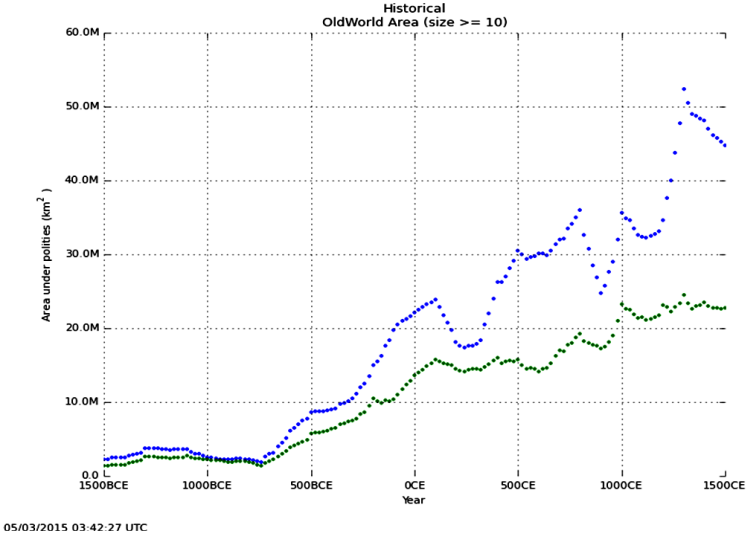

Figure 2. Area under agrarian (green) and all (blue) empires, interpolated from the TCTG13 data of Turchin et al. (2013).

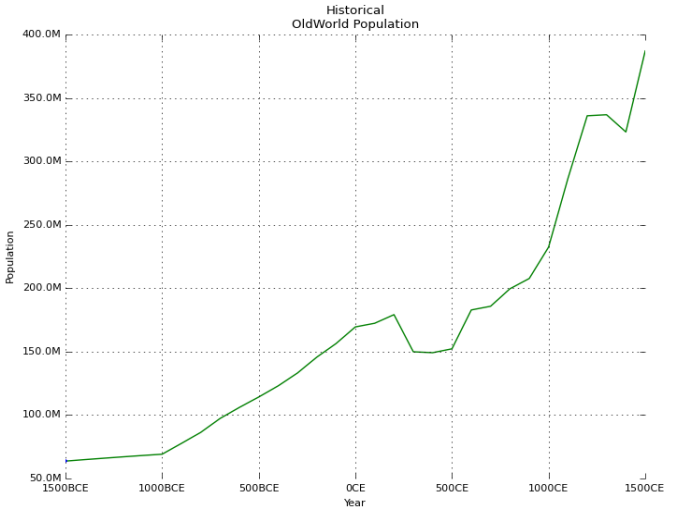

Figure 3. Old World population from the KK10 dataset of Kaplan and Krumhardt, 2008.

The Demographic Model

To predict the agrarian Old World populations produced by the ADSM model I employed the large-scale demographic model of Bennett (2015). The model assumes that a larger state, by its enhanced organization, is able to support greater crop yields, and hence carrying capacity advantage, over a smaller state. In particular, the model posits a simple, uniform threefold carrying-capacity advantage of multi-region 'empire' polities over single region 'hinterland' polities in the TCTG13 dataset between 1500 BCE and 1500 CE. The regional population P_r is grown logistically following Equation 1:

$$\dot{P_r} = \beta(\tau_r)\left(1 - \frac{P_r}{K_r}\right)P_r \quad (1)$$

where $\beta(\tau_r)$ is the net birth rate associated with the polity type τ (empire or 'hinterland', defined above) in region r and K_r is its carrying capacity.

Regional carrying capacity estimates K_r follow Equation 2, which reflects mean population densities ρ for each polity type and an exogenous, temporally varying regional agricultural intensification schedule $I_r(t)$:

$$K_r = \rho(\tau_r)\, I_r(t)\, A_r \quad (2)$$

where $\rho(\tau_r)$ is the population density associated with the polity type τ in region r and A_r is its area.

In cases where regional population exceeds carrying capacity ($P_r > K_r$), as happens when an empire collapses to hinterland, the precipitous population collapse is modeled as

$$\dot{P_r} = -(P_r - \delta K_r) \quad (3)$$

where δ is the fraction of the carrying capacity to which the population falls. Summing regional populations as their polity types change historically yielded plausible predictions of the overall demographics in the subsets of the Old World discussed above. Famines and plagues are not modeled.

Mean historical hinterland and empire population density estimates (ρ_h= 4 people/km^2, ρ_e= 12 people/km^2) and the empire net birth rate (β_e= 1.71%/year) are based on scaled population and agrarian yields from medieval England between 1150 CE and 1300 CE (Campbell, 2010). The hinterland net birth rate (β_h = 0.2%/year) is based on KK10 historical population data in regions without empires. Values for β reflect the typical net birth rates at low actual population densities for the two types of polities. The effective birth rate of Equation 1 depends on both these base rates and the residual carrying capacity available. As

a growing population saturates regional carrying capacity, the effective birth rate becomes a fraction of the base rates.

Lastly, the model implements a within-polity migration scheme in which a fixed fraction (μ) of a polity's population moves between its regions each time step according to available residual carrying capacity. Fraction values for population migration (μ= 0.2%) and death from carrying-capacity collapse (δ= 80%) were estimated from overall fit to TCTG13 historic polity and KK10 population data.

The agricultural intensification schedules $I_r(t)$ employed by Equation 2 reflect empirically estimated regional variations in agricultural yields, including an apparent Old World-wide doubling of intensification circa 1000 CE. Figures summarizing the geographical and temporal intensification schedules are included in Supplemental Materials.

The demographic model makes no assumptions about how polities rise and fall and thus can serve to predict the demographics that different imperiogenesis models generate. Here I employ the demographic model to drive the ADSM model, evaluating its population predictions against the historical record.

I now describe the agrarian demographic-structural model itself.

An Agrarian Demographic-Structural Model

The central premise of Goldstone's demographic-structural theory posits that, following an 'integrative' period of growth and internal cooperation, progressive immiseration of the major interacting groups of a society—its workers, elites, and state agents—eventually leads to social upheaval and sometimes rebellion (the 'disintegrative' period) after all three find themselves competing for control of a constrained pool of resources in order to pursue and enjoy their roles. The particular source and coupled sequence of immiseration is different for each group, but all are rooted in demographic growth outstripping the opportunities structurally available to each. Goldstone, and Turchin and Nefedov in turn, have outlined in some detail the sequence and consequence of these stresses on the tripartite structure of society, making predictions about several important societal characteristics, including the re-distribution of wealth among the groups and aspects of intra-elite competition that can precipitate rebellions.

Demographics and Structure

I approximate the details of this literature for agrarian polities during the period of investigation. In particular, my agrarian demographic-structural model is a spatially-explicit geopolitical agent-based model in which agrarian polities are composed of only two political-economic sectors: a group of farmers (‘peasants’) working the fields to produce the food to support both themselves and a smaller group of warriors ('elites') who provide protection against banditry and

annexation, as well as manage and invest the infrastructure (here people) required for production of food. Thus each sector, involving specialized labor, provides a valuable public good for the other. The two-sector cooperative involves a social bargain, however explicit and at least initially felicitous, that, in exchange for food, the large-scale stability provided by the elites increases the total number of people in both sectors compared to a 'hinterland' polity where the roles and organization may not be as specialized (and hence not as effective).

No doubt the elite organization has some hierarchical political structure but I do not model it explicitly. The elites here are best thought of as the combined entourage of some number of aristocrats and their military retinue who are able to cooperatively provide the protective and expansive function outlined below. Similarly, farmers have no modeled organizational structure. Further, there is no promotion or demotion of individuals between sectors; people are born to their station and remain there.

In the model, farmers play a relatively minor role; as land becomes available through the agency of their associated elites, farmers migrate and grow into that opportunity and thus support themselves and the elites. The actions of the elites, on the other hand, who perform and bear the costs of the defense and annexation drives the overall behavior of the model. Their response to swelling and then dwindling opportunities (the demographic-structural crisis) precipitates political upheaval and regime change. Indeed, the ADSM model focuses exclusively on demographic-structural factors and mechanisms that impact the behavior of elites, not the peasants.

As in the Bennett (2015) demographic model, there are two types of polities: Single-region 'hinterland' polities and multi-region 'empires'. Each region's population is split between the sectors: farmers and elites. Each region's total carrying capacity is divided between the sectors in a fixed fraction, ε, set here uniformly at 20% to the elites. The constant ε reflects a plausible guess of total elite fraction from historical estimates of medieval and ancient societal structures (see Turchin and Nefedov 2009, for example). The particular value of ε serves largely to scale the size of projected armies (and their deaths). I adopted a uniform value for simplicity.

The population of each sector grows separately and logistically per region, following Equation 1:

(4) $$\dot{F}_r = \beta_F(\tau_r)\left(1 - \frac{F_r}{(1-\varepsilon)K_r}\right)F_r$$

(5) $$\dot{E_r} = \beta_E(\tau_r)\left(1 - \frac{E_r}{\epsilon\, K_r}\right)E_r$$

where F_r and E_r are the current populations of a region's farmers and elites and $P_r = F_r + E_r$. Although a good case has been made that the elites would enjoy a higher net birth rate than the farmers, for simplicity I assume they are identical ($\beta_F(\tau_r) = \beta_E(\tau_r)$) and, following the demographic model, are equal to either β_e or β_h depending on the polity type. In the case of carrying capacity collapse, Equation 3 applies to each sector separately using the same collapse fraction δ.

Although most analyses of the demographic-structure theory track and predict various measures of wealth and its (mal-)distribution between sectors, my model does not. I found it sufficient to track *residual opportunity* for the two sectors in order to drive their behavior. The residual opportunity for farmers and elites is defined by the $\left(1 - \frac{F_r}{(1-\varepsilon)K_r}\right)$ and $\left(1 - \frac{E_r}{\epsilon K_r}\right)$ terms of Equations 4 and 5, respectively. The model begins by seeding an initial set of empires based on the TCTG13 historical data at the start of the period in the landscape of interest. Following Bennett (2015), the initial KK10 historical population is divided regionally according to polity type (hinterland or starting empires), and each region's population is then divided into farmers and elites according to ε. The simulation then unfolds independently of the historical record according to the behaviors below. In all the simulations in this paper, the time step was 2 years.

Armies and Annexation

As the number of elites increases, their ability to project militarily their power increases and they become able to annex neighboring agricultural regions. However, power projection is subject to its own (spatial) logistical constraints as it takes more people in the elite entourage to support an army prosecuting a battle in a distant neighboring territory. In particular, I follow both Artzrouni and Turchin by assuming that all the polity's elites are effectively concentrated into a central depot and then spread out to various political borders. The force, H, that can be projected to a contested region is given by:

(6) $$H_{polity}(d) = \left(\frac{\sum_{polity} E_r}{\left|Borders_{polity}\right|}\right) L(d)$$

(7) $$L(d) = \frac{1}{\left(1 + \frac{1}{s}\right)^d}$$

where $H_{polity}(d)$ is the size of the polity's army at the region d kilometers away from the polity's depot, $\sum_{polity} E_r$ is the elite population of the entire polity and $\left|Borders_{polity}\right|$ is the number of polity regions that have *political* borders requiring defense or a battle to annex. Unlike Artzrouni, I assume *physical* border

regions such as deserts and impenetrable mountains require a negligible defensive army. Sea-based attacks between littoral border cells are discussed below.

The function $L(d)$ encodes how projected military power attenuates with distance. As supply lines lengthen from the depot to the contested border, some part of the army must support the soldiers doing the fighting. As distance increases, those support troops themselves need support, etc., leading to the recursive form of $L(d)$ in Equation 7 above. Boulding's (1962) derivation of this relation is given in Supplemental Materials. The key logistic scaling parameter, s, is the number of fighting soldiers supported per supply soldier per 100 km; larger s implies a larger projected force at a given distance. In Supplemental Materials, I employ the TCTG13 historical data to estimate s to be 1.5 before 700 BCE, 1.95 thereafter in this period.

At each time step every polity decides whether to launch offensive attacks on some of its neighbors. In the current model, one of the immediate neighbors of each political border region of a polity is selected randomly. If the chosen region is both agricultural and occupied by a different polity an 'order' to attack that region is written. All the orders from all the polities, including hinterland regions, are collected, randomly shuffled, and then executed. The stochastic selection of at most a single neighbor per border region effectively limits the total number of attack orders written per polity each time step; shuffling the orders randomly ensures there is no overall spatial bias to empire growth.

As each order is executed it is first checked that the offensive polity still occupies the attacking region (they could have been attacked themselves and annexed) and that the attacked region is still occupied by the same 'other' polity when the order was written. If not, the order is rescinded. Otherwise the power of the polity armies at the attacked region is assessed using Equation 6 and if the ratio of the offensive to defensive power is greater than a threshold, Δ, the attacked region is annexed into the victorious polity.

$$\frac{H_{attacker}(d_{attack})}{H_{defender}(d_{defend})} \geq \Delta \tag{8}$$

where d_{attack} is the distance from the attacker's depot to the chosen region and d_{defend} is the distance from the defender's depot to the same region. I set Δ to 1.2, requiring at least a 20% military advantage to be victorious. Larger values of Δ require larger elite populations (for a given s) to achieve the same territorial expansion, which interacts with elite immiseration described below.

My war and annexation scheme is similar to those described by Collins (2010), Turchin (2003), and others, but has some important differences. First, the determination of the polity's single depot is chosen as its geographic center at the

formation of the polity. Turchin employed a similar rule, but permitted the central depot to move to a new geographic center as territories were annexed. I experimented with both schemes but settled on fixed depots for simplicity. Both schemes lead to polities that radially project power from a central location and hence have a more or less circular shape, geography permitting. Further, both approaches tend to expose one or more flanks to weakening support as the stochastic choice of territories exceeds the power projection permitted by *s*. In the case of a movable depot the polity (of a relatively fixed size) can move across the landscape that the fixed depot scheme does not permit. However, my experiments and analysis of the actual historical data suggest that fixed depots are more common and somewhat in keeping with maintaining 'sacred' locations associated with states (Atran et al. 2007).

A second difference is that there is no explicit 'asabiya' term reflecting any increase in the military effectiveness of either army because of heightened camaraderie or dedication to a cause, for example. While military history is replete with examples of this effect and its often-decisive impact on battles (Turchin, 2006c), I elected to keep the model simple and avoided modeling that effect during battles. Further, unlike the A2 model, there is neither enhanced asabiya in border regions, nor eroded asabiya in core regions; a polity's elites are assumed to be uniformly enthusiastic in their pursuit of annexation and defense. Nevertheless, increased asabiya does make an appearance during collapse and civil war, discussed below.

There is little cost/benefit computation by a polity about *which* battles to prosecute; they are instead stochastically chosen. A polity has no preference for, say, regions at a similar latitude (Turchin et al. 2006a), or for reclaiming 'sacred' regions that were annexed by a neighbor. Further, the current model evenly divides the 'depot' force among all the political borders. There is no provision for concentrating forces on a particular region or enemy polity. These and other advanced military strategies and tactics are not modeled here as, once again, I assume that on average the effect of these will be to pursue all possible opportunities with the force at their disposal.

The only preference employed by the model involves the expected productivity of the targeted agricultural region. Since the population density values are fixed, expected productivity is reflected by the regional intensification factor $I_r(t)$ in Equation 2. If the intensification factor of the targeted region is below a threshold, a, the region is deemed 'not worth it' and the order is not written. Throughout all simulations the threshold a was set to 0.6, the typical intensification lower bound from the demographic model. This simple preference, however, coupled with exogenous changes in intensification factors (described in Bennett (2015)) is critical to matching the actual historical pattern observed in Europe, as discussed below.

In their TCTG13 model, Turchin et al. (2013) augmented the power of a defensive army according to the elevation of the attacked region, making mountainous agricultural regions more difficult to annex. I pursued a different approach. Bennett (2015) declared any agricultural territory above 3 km in elevation to have a very small intensification factor, thus making it not worth the while of any polity and effectively erecting a physical barrier to annexation. No doubt higher elevations were annexed and claimed as parts of polities more for their defensive than productive provision. Indeed, after analyzing the number of regions in the TCTG13 dataset that were annexed at different elevations, I found that regions below 3 km were routinely exchanged throughout the data record with a clear drop off above 3 km (Supplemental Materials). These high elevations are found only in the Himalayas and western China. Thus, unlike Artzrouni, who declared the Pyrenees and the Alps impenetrable, my model permits their annexation at no additional cost over a marchland plain, a choice also critical to the formation of Europe using my model. Of course, employing an agricultural measure as a proxy for a defensive enhancement is a confusion of concerns, however serviceable, and should be clarified in later models.

War Deaths

Elites attempt to annex territory in order to increase the total elite opportunity available to themselves and their descendants. But a newly annexed territory is already occupied by some elites—of the other polity. Thus, when a battle is successfully waged, opportunities must be made and some fraction of both the victorious and the vanquished *projected* elite armies at the battle site are killed.

The elite army is formed by contributions (here assumed to be 100%) from all the polity's elites in all regions, which are then redistributed to the contested borders. Therefore any battle deaths should be made proportionally to elite population distribution over *all* the polity's regions, not just within the contested region. In this way battles, especially unsuccessful ones, can slowly weaken the entire polity and eventually its ability to support its borders, even causing runaway loss of all territories.

The two death fractions, $\delta_{victorious}$ and $\delta_{vanquished}$, are set at 10% and 80% respectively. I arrived at these values empirically by varying these parameters leaving all else fixed and observing that the overall predicted demographics of the world remained plausible. In particular, if $\delta_{vanquished}$ is not very high, vanquished polities remain strong and delay annexation, and the predicted world population overshoots the observed population; $\delta_{vanquished}$ needs to be above 50%, but much beyond 90% the predicted population drops well below the observed.

In the current model, there are no deaths of farmers on either side during battles, nor is there loss of agricultural productivity. While it would be straightforward to add more death fractions for the farmers similar to the elites

(except they would apply locally), I elected not to complicate the model and assume, instead, that the vanquished farmers continue to work under the victorious elites.

Immiseration, Collapse, and Civil War

Territorial expansion continues until the spatial logistic limit set by s (and Δ) is reached and the growth of elites is constrained both spatially and by carrying capacity limits in the occupied regions. As these limits to growth are approached, the elites become immiserated, cease their cooperative behavior, and engage in intra-elite competition. The farmer population, of course, has been growing as well and their opportunities have also been shrinking. The demographic-structural theory predicts many consequences of this farmer immiseration such as real wage loss, reduced health as diets become poorer, rural flight to urban locations, and food riots. I model none of this and assume that in spite of the elites own increasing immiseration they retain enough cohesion to suppress any insurrections by the farmers. Thus, farmers come to saturate their fraction of the carrying capacity, but their fate remains controlled by the elites.

I model elite immiseration as a simple threshold, Θ, on a region's consumed elite opportunity $\left(1 - \frac{E_r}{\epsilon K_r}\right)$ and set Θ to 90%. Thus when only 10% of a region's elite opportunity remains, that region's elites become disgruntled and launch a rebellion. Goldstone (1991) suggests that some (large) fraction of a polity's regions must become miserable before a rebellion occurs, but I model the trigger as a single region. Empirically this turns out to be a plausible assumption, since most regions are annexed around the same time and grow at the same rate so most regions are not far behind one another in saturating their opportunities.

In my model, once a rebellion is raised, the polity collapses into two or more factions. The specific number of factions depends on different experimental probability distributions. I investigated several distributions of up to five factions with decreasing probability (see Supplemental Materials). Increasing the number of factions appeared to make little difference to the predictions at this scale as long as the chance for two factions dominated the distribution. I settled on a simple fission into two factions for these simulations.

To form these factions, two (or more) of the collapsing polity's border regions are chosen at random and assigned to new faction polities. Then, in a round-robin fashion, each faction chooses a region on its expanding border, regardless of its power, that was part of the original polity but has not yet been assigned to another faction. In addition, there is a chance, f, that a faction could skip its turn. This has the effect of creating unequal-sized factions, whose resultant power then differs during the ensuring civil war. Throughout the simulations, f was set to 50%, an even chance of not acquiring a region per round. Of course, a faction occupying a single region is converted to hinterland.

Reviewing the historical data in TCTG13, portions of empires sometimes collapse below the minimum polity scale and appear to revert to hinterland. Although a few of these historical regions likely collapsed to smaller states, I assume that all fall into relative disarray and are unable to form and maintain a cohesive polity. I found all episodes in TCTG13 where an empire lost at least 10% of its prior century's territory and then computed the fraction of the actual loss (often much higher) that went to apparent 'hinterland' (as opposed to other large states). The odds that more than 50% of lost territory collapsed to hinterland are about 25% (76 out of 294 loss events). Based on this value, I set the odds (f_h) for a formed faction to fall immediately to hinterland to 25%. Of course falling to hinterland causes substantial population loss for both sectors and eases the field for surrounding states, including new factions, to rapidly annex this territory.

Factions are formed because of a rebellion by (at least) one immiserated region. The split into factions does nothing to reduce the misery of that region and any faction that inherits it would itself immediately collapse, recursively, until that region was reduced to hinterland. To avoid this either elite opportunity or their tolerance for misery must temporarily increase; I model a mixture of both. I chose a 'grace period' (g) of 25 years, roughly one generation, in which the factions might stabilize into longer-lived polities (or fall again). I achieved this both by eliminating a fraction of elites from each faction (assumed to perish during the spasm of collapse) and by raising the misery threshold for all factions for that period of time, in some combination controlled by a mixing parameter, *c*. The increased misery threshold reflects a temporarily increased asabiya by *all* factions to their new respective causes.

To estimate the new misery threshold, Θ′, given *c* and g, I first solve Equation 1 for the predicted population *P* after a time interval *t*

(9) $$P = \lambda_P(P_0, \beta, K, t) = \frac{KPe^{\beta t}}{(K + P_0 e^{\beta t - 1})}$$

where P_0 is the starting population, K is the carrying capacity, and β is the base birth rate. I employ Equation 9 to estimate Θ′as $\lambda_P(\Theta, \beta_e, 1, g \cdot (1 - c))$ using fractions of carrying capacity and populations for K and P_0. Similarly I solve Equation 9 for P_0

(10) $$P_0 = \lambda_{P_0}(P, \beta, K, t) = \frac{KP}{(Ke^{\beta t} + Pe^{\beta t - 1})}$$

to compute the fraction of elite carrying capacity, hence population, such that it would take g years to attain normal misery $\lambda_{P_0}(\Theta, \beta_e, 1, g \cdot c)$. I chose c to be 80% death, 20% increased tolerance, implying a Θ′ of 91% (a 1% increase in

tolerance) and a drop in the elite population to 86% of εK. After adjusting the elite population, a timer is set on each factional polity for the grace period; once it expires, the polity's misery threshold is reset to Θ, reflecting a loss of heightened support to the faction and perhaps a desire to return to normal life. Of course, any regions that become disgruntled at that point will trigger a collapse of the polity.

Once cohesive factions are formed they operate as usual and can prosecute wars, often against their opposing factions, since by construction they share an extensive political border. These civil wars are fought identically to normal annexation wars between different polities.

When a faction is formed, its depot is assigned to the region closest to its geographical center. By construction, starting from the border and working inward, the faction depots are always distal to the original polity's depot. Thus, assuming that either or both factions survive, they annex territories both toward the original center (the civil war) and away from it, into relative frontier. This factional expansion is responsible for most of the model's large-scale diffusion of empires throughout the Old World.

When factions form, their relative projection power roughly halves (their population splits and their spatial extent halves). Certain regions under each faction's sway may no longer be well defended with these reduced means. In addition to any surrounding polities, hinterland regions next to these vulnerable regions may then be able to annex them and thus arise to become small polities themselves. These new polities then enjoy the improved carrying capacity and net birth rate of empires and may grow by annexing additional hinterland and weakened faction regions, similar to the faction itself. Indeed, about 3% of large polities are created in this indirect fashion during simulations of the Old World. This expansion behavior resembles the 'reflux' behavior that Turchin observed in his A2 model, described earlier. However, in the ADSM model polities do not form spontaneously from hinterland regions in response to meta-ethnic border threats, but instead form in response to power shifts during polity collapse and weakening.

Table 1 summarizes the parameters of the ADSM model, exclusive of the parameters from the Bennett (2015) demographic model. Except for s, the logistic supply parameter, and f_h, the chance of a faction collapsing to hinterland, both of which are estimated from the TCTG13 dataset, all other parameter values are plausible estimations that permit the ADSM model to fit the historical data reasonably. They stand as predictions that may be verified or improved using other, independent historical datasets. For example, military historians may be able to suggest better mean values and perhaps ranges for the 'war' parameters (s, Δ, $\delta_{victorious}$, and $\delta_{vanquished}$) over this period.

Table 1. ADSM model parameters

Parameter	Value
Elite carrying capacity fraction (ε)	20%
'Worthwhile' agricultural productivity (a)	0.6
Soldiers per supply soldier per 100 km (s)	1.5 before 700 BCE, 1.95 thereafter
Victory power ratio (Δ)	1.2
Elite death fraction ($\delta_{victorious}$)	10%
Elite death fraction ($\delta_{vanquished}$)	80%
Elite immiseration threshold (Θ)	90%
Grace period after collapse (g)	25 years
Death/tolerance mixing parameter (c)	80% death
Number of factions per collapse	2
Chance to pass acquiring a region into faction (f)	50%
Chance of a faction collapsing to hinterland (f_h)	25%

I summarize the model's sensitivity to parameter variation below. Regardless of their settings, however, since the parameter values apply to *all* polities, except for the stochastic choice of which regions to attack and how to form new factions, there is no variation in polity response to the opportunities and challenges each polity faces.

Further, the values represent *mean* historical values, so idiosyncratic tendencies and behaviors, such as variations due to culture or political structure, play no role in the ADSM model; these variations are assumed to have little cumulative impact (which is likely false, as future empirical analysis may demonstrate). Nevertheless, as will be seen, the impact of stochastic variation, especially when coupled with geographical and agricultural constraints, can be substantial.

Behavior and Results

In this section I first review the model's general behavior in a simple abstract environment. This is followed by a comparison of the model's overall behavior with the Old World historical record, a summary of its sensitivity of that behavior to certain parameter variations, and then a discussion of the model's 'secular cycle' longevity variance with historical polity duration. Later sections discuss the model's behavior in the Asia, Europe+ME, and India+ME theaters.

The spatial and temporal behavior of the ADSM model in different theaters can be seen in several accompanying animations (see Supplemental Materials). In these movies, a different colored marker is randomly assigned to each polity to distinguish it for its duration; markers are reused. All movies employ the

parameters in Table 1 and depict a single run of the model; composite statistics over 20 runs are presented below.

General Behavior

One movie (SM_Center) shows the model's general behavior in a simplified, completely agricultural, rectangular landscape with no elevation gain and no oceans. An initial empire in the middle of the region seeds the simulation. There are three broad phases of activity (see Supplemental Materials). First, the seed empire expands relatively rapidly as its population grows, funding an increased army that is projected into the surrounding hinterland plain. Territorial growth slows and eventually the power of the hinterland balances this power and the empire ceases to expand. A few years later immiseration sets in and the empire collapses into two factions, both of which reclaim territories and expand away from the original center. At this point there is also a flurry of state formation as some relatively more-powerful hinterland regions annex the peripheral regions of these weakened factions. Thus, step-by-step, the space is filled with polities that are a consequence of these initial collapse factions. Once the space is saturated, behavior settles into a third phase of fluctuating exchange between a few large empires, all with similar sizes dictated by *s*. The slow change of borders in the last phase is due largely to the stochastic order of annexations, a subtle but important effect. Rapid change of borders occurs if a large faction happens to fall to hinterland. Overall, this spatial behavior resembles Turchin's A2 model, but is driven by different underlying mechanisms.

I now turn to the model's behavior and overall predictions in the various subsets of the Old World.

The Old World

The statistical results from 20 simulations of ADSM of the Old World are shown in Figure 4 below. Each simulation begins in 1500 BCE with three seed empires: Thutmosid Egypt and the Hittite empires in the Middle East and the Shang Dynasty in Asia. The expansion fate of the western and eastern empire centers is very different and I discuss their detailed behavior below. Overall, the total area occupied by agrarian empires expands at a rate that is very similar to the observed, historical rate. This expansion rate emerges from the underlying demographic-structural behavior of the model. The early historical expansion delay before 600 BCE is only partially matched in the model, due to geography (see below) and the lower value of *s* in this early period. Moreover, the model does not predict the dramatic expansion of empire area during the rise of Rome and the Han around 0 CE. While the population prediction tracks the historical population, if biased higher, the underlying demographic model does not predict the Justinian plagues and world famines or the effects of their rebound and, thus,

the exogenous doubling of intensification by the model at 1000 CE overshoots the historic population. The model produces a comparable but slightly larger number of long-lived empires as the historic record.

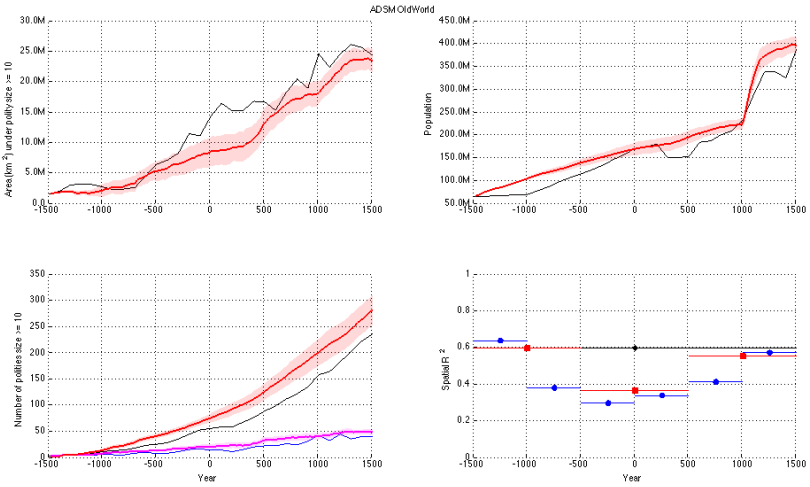

Figure 4. Predictions for the Old World. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. See the SM_OldWorld animation.

The occupied area graphs, coupled with the graph of number of empires, both on the left side of the figure, provide an indication of when empires arose but not whether they arose in the same region as the historical data. Following Turchin et al. (2013) I evaluated this aspect of model behavior using a spatial R^2 statistic. The results are shown in the lower right panel of Figure 4. As the simulation unfolds the system counts the number of times each region is occupied by a large empire (≥ 10 regions) for the trailing 500 (blue dot interval), 1000 (red square interval), and 3000 (black diamond interval) years. The system performs this count for both the historical and predicted empire locations and computes the R^2 metric for each interval of interest as they occur. Turchin et al. (2013) reported only the trailing 1000-year (red) metric.

Although by construction each simulation starts by matching the world precisely, the Old World simulation spatially diverges rapidly from the actual expansion, as noted above when discussing the initial premature area expansion. This is reflected in the rapidly declining values of the 500-year metrics shown in Figure 4. Nevertheless, on the 1000-year (and 3000-year) scale the simulation does reasonably well. Overall on the 3000-year scale the model explains 60% of the spatial variance of the historical data, slightly lower than Turchin et al. (2013).

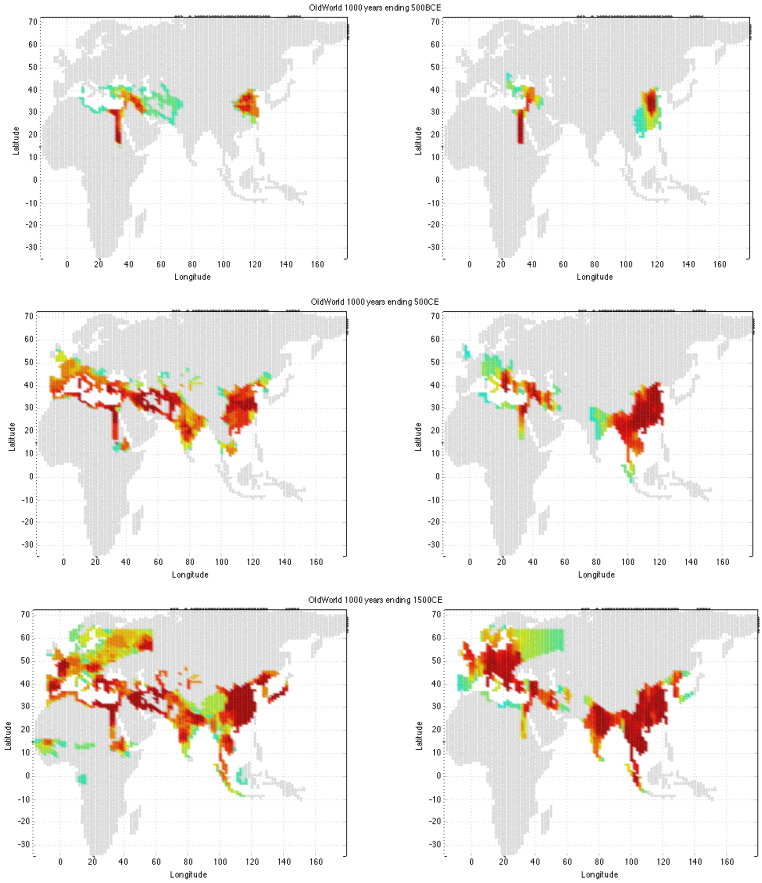

Figure 5. Imperial density for the Old World. Left panels reflect historical data from the TCTG13 dataset for large-scale agrarian states (≥ 10 regions); right panels reflect model predictions. Red (cyan) indicates more (less) frequently occupied regions.

The above figure compares the historical and predicted large-scale imperial density at the end of each elapsed millennium. The left panels reflect historical imperial densities; right panels are model predictions. Red indicates regions more frequently occupied by agrarian empires; cyan indicates regions least frequently occupied by agrarian empires

Sensitivity of ADSM Behavior in the Old World to Parameter Value Variation

I assessed quantitatively how well the model fits the historical data by combining the millennial (1000-year) R^2spatial metric with estimates of millennial mean area errors evaluated in the Old World theater. The details of the sensitivity analyses are available in Supplemental Materials. Overall, like the TCTG13 model, the ADSM model using the nominal parameter values in Table 1 is able to explain around two-thirds of the overall and millennial R^2spatial variance in polity density while also predicting empire area values often within 20% of their historical values.

Model predictions depend sensitively on values for the parameters controlling war and annexation. Increasing or decreasing the logistical power projection parameter (s) increases or decreases, respectively, the number, size, and, hence, population of large polities, lengthening or hastening their time to collapse. Increasing the projected power ratio (Δ) required to annex territory also severely limits the size and spread of polities. The ability of the model to predict the historical record in these cases is substantially compromised. Finally, the death fraction for the vanquished (but not the victorious) population is also critical to proper fit, since it controls the size of the annexed elite population, hence residual opportunity, lengthening the time to collapse as that value increases.

Somewhat surprisingly, the model predictions do not depend as sensitively on precise values for most of the parameters that control the polity collapse mechanism. As long as the probability of collapse to at least two factions (f) is above 40%, there is no significant impact on the model's results for most collapse parameters. The exception is the residual opportunity tolerance threshold (Θ). If it is increased to 100%, none of the initial polities collapse and no spread of agriculture occurs. Lowering it to 50%, however, while having only a modest impact on the area occupied by large-scale empires, substantially increased the total number of smaller polities generated.

Secular Cycles in the Old World

The secular cycles discussed by Turchin and Nefedov (2009) typically involve historical polities, such as the Roman Empire, undergoing repeated demographic-structural crises, sometimes with civil wars. The ADSM model has no notion of when a rebellious faction persuasively claims the mantle of its parent polity, thus

providing symbolic or political continuity (or reformation) after a civil war. This accounting difference is one of the reasons why the ADSM model produces an apparent excess of polities and consistently overestimates the cumulative number of empires.

Nevertheless, the rise and collapse episodes of polities generated by the ADSM model are a direct implementation of demographic-structural secular cycles. Figure 6 shows the normalized distribution and variance of collapsed large-scale ADSM polity duration sampled every 25 years over 20 model runs for the Old World.

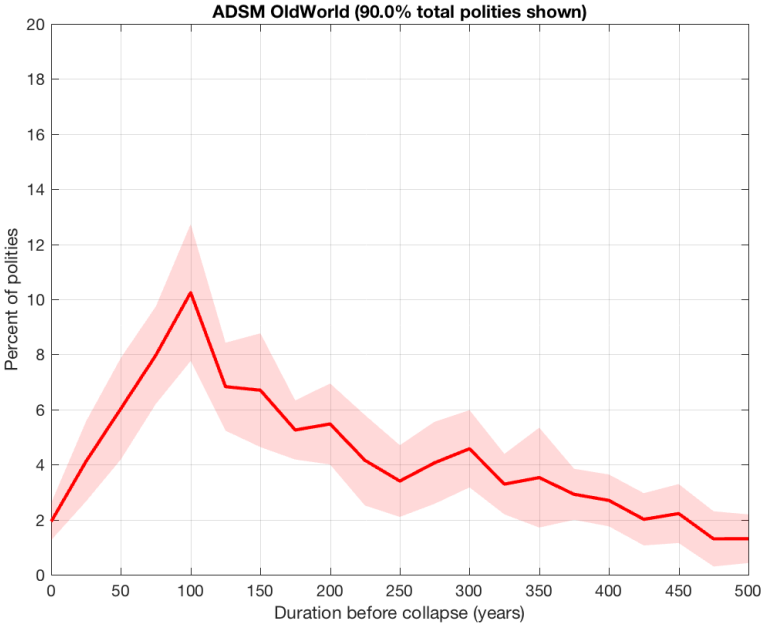

Figure 6. Normalized distribution and standard deviation of ADSM secular cycle duration in the Old World sampled every 25 years.

Historical analyses of apparent demographic-structural secular cycles suggest the length of cycle should vary between 150 and 300 years (Turchin and Nefedov 2009). As can be seen, a substantial number of large-scale ADSM polities are relatively short-lived (~25% survive less than 100 years), primarily reflecting factions that were unable to stabilize against inner or outer forces with the resources available. Recall that the cumulative number of polities shown in the lower left panel of Figure 4 reflect polities with duration of at least 100 years. The

majority of these longer-lived polities (75%) have duration between 150 and 400 years.

In spite of employing empire population and carrying capacity parameters estimated from a single 150-year secular cycle (medieval England between 1150 CE and 1300 CE), the ADSM model demonstrates that the relationship between a polity and its geographical settings serve to broaden secular cycle duration. The longest-lived modeled polities survive because of repeated wars over geographically undesirable regions at the periphery of larger empires. These wars temporarily reduce their elite population, thus averting collapse, yet avoiding annexation before the attacking polities collapse themselves.

The TCTG13 dataset reports large-scale polities that survive for at least a century. Figure 7 compares the normalized distribution and variance of ADSM polity duration with the normalized distribution from that dataset at the same time resolution (century).

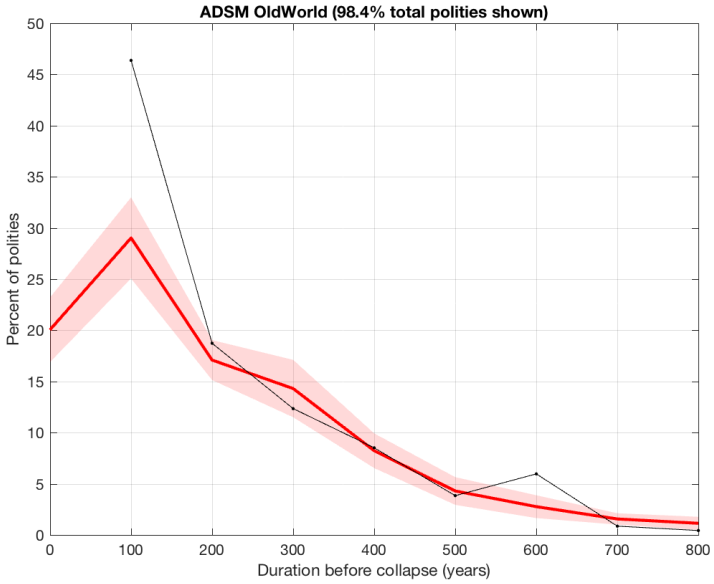

Figure 7. Normalized distribution and standard deviation of ADSM secular cycle duration in the Old World sampled every century (red) compared with the normalized TCTG13 empire duration data (black).

In the following sections I discuss the model's behavior in the Asia, Europe+ME, and India+ME theaters, leading to two experimental investigations that improve the model's overall performance in the Old World.

Asia

In spite of the somewhat encouraging statistical results in the Old World, analysis of the different subset regions gives us pause. The results for the Asian theater, seeded by the Shang in 1500 BCE, are shown in Figure 8. The area plot makes it clear that the model expands prematurely compared with history with all the expected consequences: Many more earlier polities (rather than several much-larger empires) lead to a 100% increase in the cumulative empire statistic by 1500 CE. Indeed, the model saturates all available agrarian space by 400 CE, nearly a millennium early. Nevertheless, the population prediction parallels the observations surprisingly well until the historical famines and plagues that were the Asian counterpart of the Justinian episodes. As expected, the demographic model predicts no growth slightly after saturation is achieved. Finally, in 1000 CE, when intensification is exogenously doubled, the model overshoots the expected population by a third largely because the (un-modeled) historical famines had decreased the actual population. Spatially, however, in spite of large variation in the short time scales, on the millennial scale the model does quite well, explaining nearly 70% of the spatial variance.

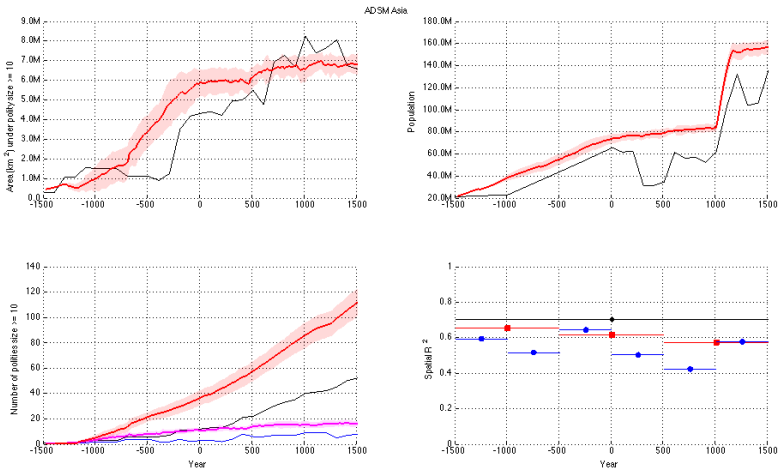

Figure 8. Predictions for Asia. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000

(red square) and 3000 (black diamond) year intervals. See the SM_Asia animation.

Europe and the Middle East

The Europe+ME simulations, shown in Figure 9, betray another historical discrepancy. The modeled offshoots of the initial Egyptian and Hittite seed polities struggle and expand very slowly with no major expansion until 300 CE, missing the earlier Roman bloom. Once the expansion begins it does so at roughly the Roman rate until it just saturates the theater by the end of the period. Historically, this predicted saturation by large-scale empires did not happen. Instead, many city-state polities underwent significant upheaval and often became smaller than the large-empire size threshold. Predicted population followed the ADSM area expansion, of course, and the intensification doubling in 1000 CE did a good job matching the population record. The spatial metric shows a good match for the first millennium, but then decreases rapidly as the Roman expansion in the second millennium is missed, finally recovering slightly by the end of the period with the expansion into the Rus; overall the model explains just over 50% of the spatial variance.

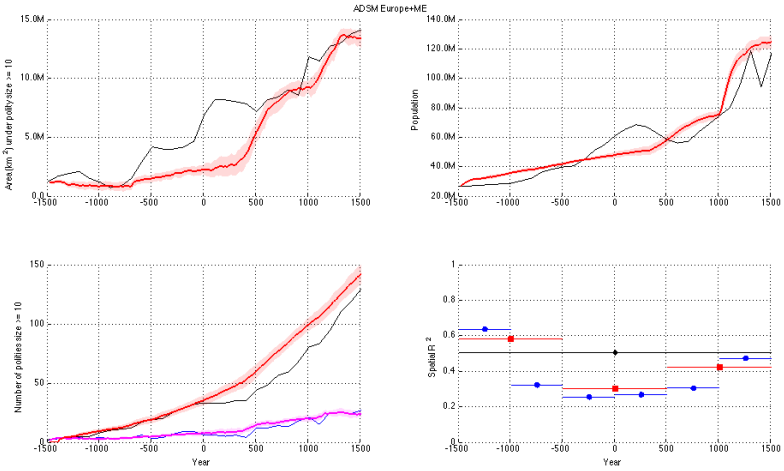

Figure 9. Predictions for Europe and the Middle East. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and

black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. See the SM_Europe+ME animation.

The striking difference between the Asia and Europe+ME simulations then lies in their early expansion behavior. While the initial Asian (Shang) empire expands rapidly, unimpeded, into the Chinese plain, the descendants of the initial empires in the Middle East have much more trouble expanding due to the small number and constraining geographical arrangement of agricultural regions available. In the absence of sea-based warfare (see below), the land-locked Middle Eastern states in the model are unable to annex sufficient regions to grow large enough armies that would permit further expansion. As a consequence, they and their successors collapse early and often, but eventually their distant children expand through Anatolia, then Greece, and finally into Italy before reaching the 'vasty fields of France' and Spain. If the Alps are asserted to be impenetrable, à la Artzrouni, by reducing the 3 km elevation limit to 1 km, the expansion is delayed still further.

The late intensification changes in 300 CE and then in 700 CE in the eastern European and then Russian plains (Bennett 2015) finally make these regions worthwhile and empires pursue their final expansion to saturate the theater. If these late intensification changes occur earlier, the model also performs much worse since it first expands northward into Eastern Europe before turning west, unlike the historic expansion along the Mediterranean first, then north into Gaul and Iberia, and finally back to the east and north into Russia. In the case of Europe, the geographic and productivity boundary conditions very clearly shape the expansion and timing of agrarian empires (Diamond 1997).

Two critical geographical regions control the expansion dynamics of early Egyptian empires in this model. The first is in the southern Nile, which is cut off from the eastern agricultural regions of sub-Saharan Africa by a single region in this encoding. If that connection were established, the model gives rise, eventually, to agrarian empires in Africa, expanding from the east rather than the west as actually happened. In the modeled geography, however, *no* African agrarian empire ever arises. Second, a few critical agrarian regions in the Levant connect the Middle East with Europe and permit empire expansion into Persia. Without those regions the east and west of Eurasia would have no influence on one another and the rise of Mediterranean Europe would take even longer.

India and the Middle East

Taken together, the Old World simulation's apparently reasonable statistical behavior then is due to the early first millennium expansion in Asia offsetting the slower-than-expected expansion in Europe in the second millennium. This observation is reinforced by the sad results of the model in India+ME.

Egypt and the Hittites in the west seed India+ME, as they did in Europe+ME. However, as noted above, the model's empire expansion there is sporadic, slow, and largely still further to the west, explaining the complete lack of predicted empires in India; no Persian empire ever arises let alone a Mauryan or Satavahanan polity. Indeed, reviewing the animation of ADSM in Asia (or the Old World), the rapid expansion indicated in Figure 8 is due to India being invaded from the east, out of Asia, rather than through Persia from the west.

To address some of these shortcomings, I performed two experiments, the first in India+ME and the second in Europe+ME. Both experiments considered whether some other boundary conditions in the historical and geographic data could account for the model's poor performance.

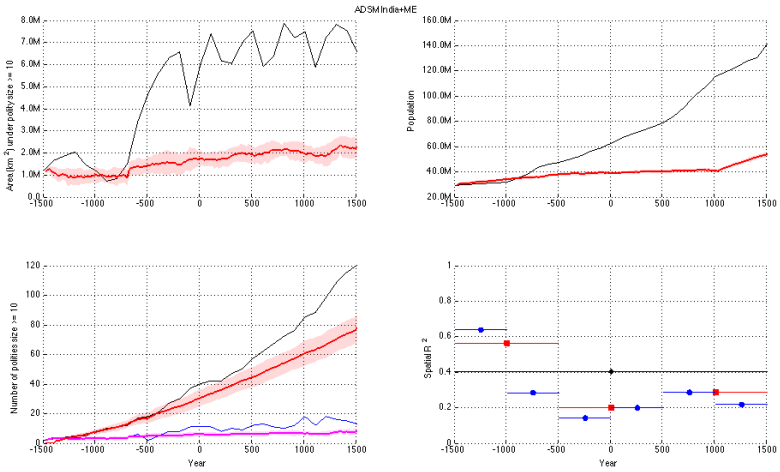

Figure 10. Predictions for India and the Middle East. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. See the SM_India+ME animation.

1. The Indus Valley Experiment

The TCTG13 historic empire data begin in 1500 BCE, just after the final fall of the Indus Valley civilization (IV) on the western shore of the Indian subcontinent. I investigated the impact of an IV empire joining Egypt, the Hittites, and the Shang

as another seeding center in 1500 BCE. The results of this experiment, shown in Figure 11, substantially improve the behavior in India+ME with only minor impact on the Old World statistics (Figure 12), as the expanding Indian empires check the expansion from Asia around 300 CE. See Bennett (2015) for a discussion of the population record and predictions in this theater.

While adding an IV civilization increases the regional fidelity of the model in the subcontinent (and increases the Old World spatial variance explained to 66%), it does nothing to accelerate the slow development in the European theater. Clearly, some other influence advanced the rapid rise of a state like Rome.

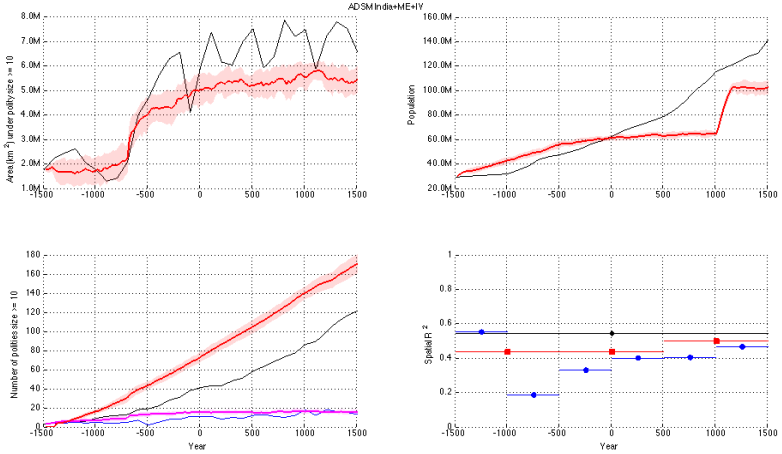

Figure 11. Predictions for India+ME with the addition of an Indus Valley civilization in 1500 BCE. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. See the SM_India+ME+IV animation.

2. Impact of Sea-borne Warfare in the Mediterranean

Following a suggestion by Walter Scheidel, I investigated the impact of sea-borne warfare on ADSM empire formation in the Mediterranean, in addition to the strictly land-based warfare explored in the previous simulations. As in the TCTG13 model, when computing possible annexations for ADSM polities, if a

littoral border region of a polity selects a sea region, a distant littoral region within a given distance (σ) is randomly selected. Assuming the selected region is agricultural and belongs to a different polity, an order is written against it and pursued in the usual fashion. Like Turchin et al. (2013), the power projection calculations employ the direct distance from the depot to the attacked region. This scheme thus increases the geographic (and hence political) connections for littoral regions.

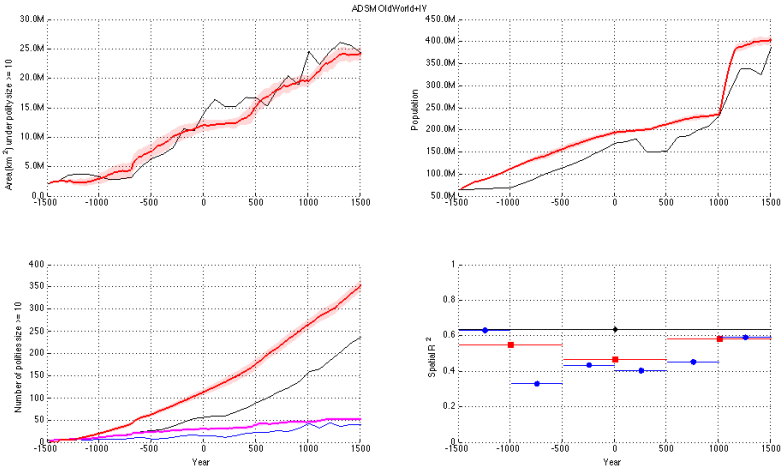

Figure 12. Predictions for the Old World with the addition of an Indus Valley civilization in 1500 BCE. See the SM_OldWorld+IV animation. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. Compare to Figure 4.

Although Turchin et al. (2013) gradually increased σ from roughly 223 km in 1500 BCE to just over 1000 km by 1500 CE, for this experiment I fixed σ at 670 km, roughly the distance from Carthage to Rome. (In all the experiments reported above, σ was set to 0 km, disabling sea-borne warfare and restricting annexations to immediate cardinal, land-based neighbors.) Further, I limited sea-based battles to the Mediterranean. While pirates no doubt plundered all the seas, for example, near the Malacca Straits or along the eastern African and Asian seaboards, my

interest is in states that employed sea power to annex and secure agrarian territory under their sway. These thalassocratic polities appear to be concentrated in the Mediterranean during this period.

Enabling sea-borne warfare substantially improves the expansion timing in the European theater; see Figure 13 and compare with Figure 9. The ability to annex distant littoral agricultural regions around the Mediterranean and establish outposts dramatically increases the population and power of early states.

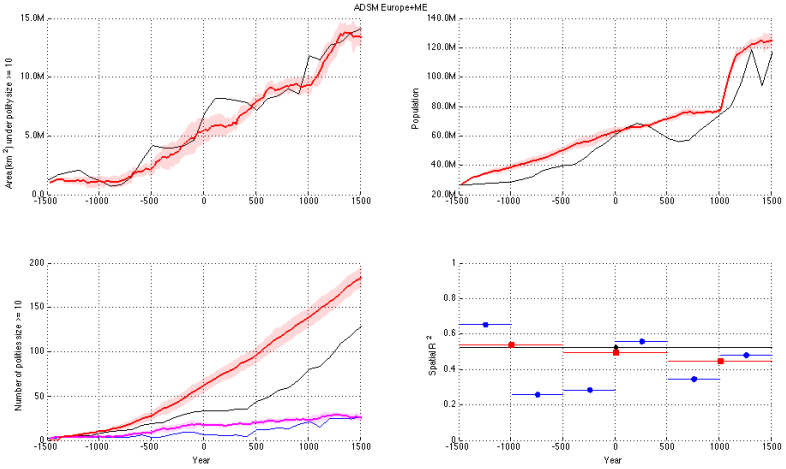

Figure 13. Predictions for Europe+ME including sea-borne warfare in the Mediterranean. See the SM_Europe+ME+sea animation. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. Compare to Figure 9.

Coupling this behavior in the Old World with the IV experiment improves the overall area and R^2 metrics, especially the 1000-year values, but at the cost of increased population and number of polities created; see Figure 14 and compare to Figure 12. Sea-borne warfare advances the model's expansion into Europe, which joins the premature expansion in Asia noted above, leading to early

increased population (and area saturation), more opportunities for collapse and hence the number of polities created.

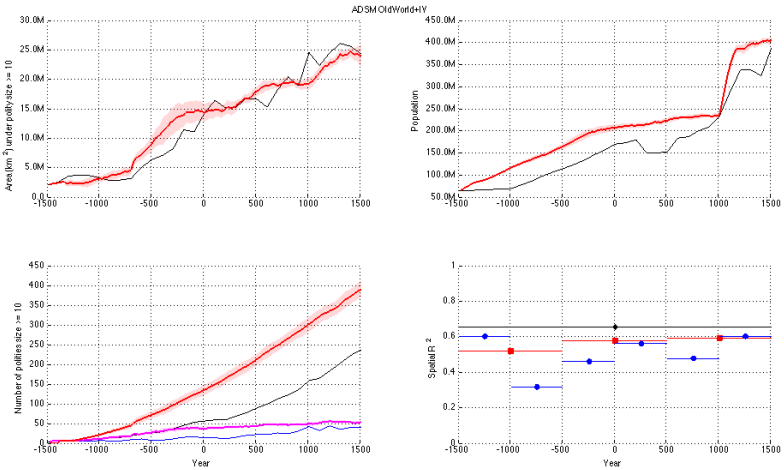

Figure 14. Predictions for the Old World including sea-borne warfare in the Mediterranean and an Indus Valley civilization. Upper panels: Predicted mean and standard deviation (red) of area under empire (left) and total population (right) every 20 years compared with historical values per century from the TCTG13 and KK10 datasets (black). Lower left: Predicted mean and standard deviation of per-century (magenta) and cumulative (red) number of large-scale states surviving a century compared with historical values (blue and black, respectively) from the TCTG13 dataset. Lower right: Spatial R^2 metrics for trailing 500 (blue dot), 1000 (red square) and 3000 (black diamond) year intervals. See the SM_OldWorld+IV+sea animation. Compare to Figure 12.

Discussion

Overall, the ADSM model supports the demographic-structural theory as a causal explanation for the three millennia of Old World agrarian history prior to 1500 CE. By coupling mechanisms of factional empire collapse and the misery-mitigating effects of population migration to previous ideas on war and annexation, the model is able to make good predictions of various large-scale historical metrics and account for much of the spread of agrarian empires in the Old World.

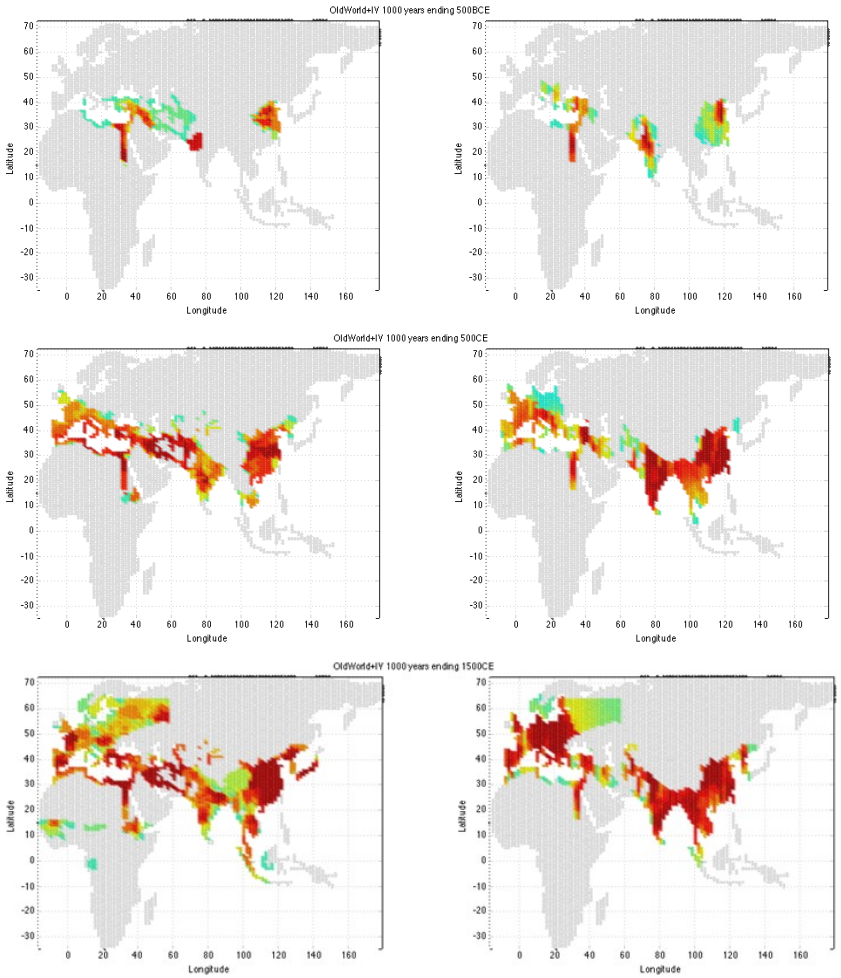

Figure 15. Imperial density for the Old World including sea-borne warfare in the Mediterranean and an Indus Valley civilization. Left panels reflect historical data from the TCTG13 dataset for large-scale agrarian states (≥ 10 regions); right panels reflect model predictions. Red (cyan) indicates more (less) frequently occupied regions. Compare to Figure 5.

The ADSM model captures many important aspects of the demographic-structural conjecture: The two-sector economies of each polity grow by extending their territories and migration of both elites and farmers ensure those acquired opportunities are exploited. However, the structural constraints of limited power

projection and capped carrying capacity are tested by an inexorable birth rate, bickering increases over dwindling elite opportunities, and finally polity collapse occurs. Factions of the original polity are formed and given spatial extent and a period of time in which to solidify their new opportunities through the auspices of inter-state and civil war. Repeated bouts of collapse propel the diffusion of states throughout the Old World.

The diffusion of states is controlled both by the model's demographic-structural factors (reflected primarily in s, μ, β, and I), but also by geographic factors such as the number and location of 'worthwhile' agricultural territories available to a polity. This is most clearly seen in the slow expansion of the Egypt and Hittite seed empires through the Levant, Anatolia, and the southeastern regions of the Mediterranean. The agricultural bridge there is small, constraining, and agriculturally disadvantaged so, in the absence of sea battles, developing power and projecting it according to our model is harder than in a region such as the broad Chinese agricultural plain. As a consequence, while the size and even shape of most resulting empires appears plausible, their locations are not, suggesting additional influences that have yet to be modeled. To increase its fidelity certain shortcomings must be addressed.

Important aspects of the full demographic-structural conjecture are not captured. More nuanced models of the internal operation of polities would track wealth concentration and redistribution, promotion and demotion of people between sectors, rural and urban flight with concomitant induced famine and pestilence increases (McNeill 1984), large-scale communication of disease, etc. These effects could all have an accelerating impact on the immiseration of both sectors (Chu and Lee 1994). Further, a 'state' sector that might serve as a check on (or amplification of) the aspirations of the elites should be included, allowing the full dynamics of rebellions to be investigated.

The various historically mis-timed territorial expansions by the model suggest additional political or geographic influences at work to accelerate or retard the growth of otherwise identical agrarian empires. Investigations here into the crude intensification schedules of Bennett (2015) and the critical connections (and disconnections) in Egypt and the Levant show that the underlying biome type data are critical to directing where and how fast agrarian empires can expand. Agricultural humanity is a shaped charge, expanding within the channels that geography (and technology) present. Improved data on plant species, their different yields, and clearance and recovery rates, perhaps incorporating climatology data, will be important to refining population and location estimates. Additionally, the demographic model of Bennett (2015) should be enhanced to predict or at least respond to exogenously-imposed major plagues and famines, which have an impact on the power and stability of afflicted empires. Finally, the exogenous intensification doubling in 1000 CE should, itself, be triggered

somehow by events in the model, perhaps the saturation of available land in the different theaters.

The ADSM model proposes that most empires arose through empire fission during collapse, but remains silent on how the initial empires arose. The circumscription theory of Carneiro (1970) posits that initial empires arose in highly fertile but geographically constrained regions that enabled managing elites to easily 'corral' uncooperative peasants who might otherwise flee. Viewing empires as a beneficial cooperative of specialized labor over a less-organized arrangement provides a plausible alternative model. Both assume a prior firm commitment to agriculture, but do not suggest how that cooperative and its benefits were discovered (or overcame its costs of increased labor and eroded hinterland traditions). It might have been easier to spontaneously discover in highly fertile regions, with or without geographical constraint, requiring less labor to enjoy the benefits of cooperation. If such locations exposed the fledgling cooperative to higher incidences of threat and banditry from the periphery, this might have sharpened the elite's specialization. It may even be that the pumping action of repeated local demographic-structural collapse bouts at a smaller scale with a (rare) faction consolidating several rivals is sufficient to raise a large state. The demographic-structural conjecture is, after all, an 'internal' circumscription theory.

However, neither consolidation nor factional alliances are in the repertoire of ADSM empires, limiting their spatial extent and increasing their number. Hegemonic empires such as Rome and the various large Chinese empires, which employ these techniques at scale, are clear outliers in the empire statistics, often at least trebling the area of the largest assumed mean empire. Once again, additional socio-political forces must be at work to engage and maintain such massive organizations. Turchin (2009) proposes that pressures from steppe nomads were a key force for inducing these increased levels of social organization, especially in China. Rome, however, may have been faced with other predatory pressures (or opportunities) and its consolidation of the Mediterranean basin must be investigated further. Even if the simple model of central power projection employed by ADSM remains plausible, some other justification would be needed for intermittently increasing *s*, the logistic support factor, for these (better-organized) empires. Modeling the evolution of enhanced social organizations was discussed in Turchin and Gavrilets (2009) and was the goal of the TCTG13 modeling effort; incorporating some of their observations would be a valuable next step.

References

Atran, Scott, Robert Axelrod, and Richard Davis. 2007. "Sacred Barriers to Conflict Resolution." *Science* 317(5841): 1039–1040. doi: 10.1126/science.1144241.

Artzrouni, Marc and John Komlos. 1996. "The Formation of the European State System: a 'Predatory' Model." *Historical Methods* 29(3): 126–134. doi: 10.1080/01615440.1996.10112734.

Bennett, James. 2015. "Modeling the large-scale demographic changes of the Old World." *Cliodynamics* 6: 57–76.

Boulding, Kenneth E. 1962. *Conflict and defense: A general theory*. New York: Harper and Brothers.

Campbell, Bruce M. S. 2000. *English seigniorial agriculture 1250–1450*. Cambridge: Cambridge University Press.

Carneiro, Robert L. 1970. "A theory of the origin of the state." *Science* 169(3947): 733–738. doi: 10.1126/science.169.3947.733.

Chu, C. Y. C. and R. D. Lee. 1994. "Famine, revolt, and dynastic cycle: Population dynamics in historic China." *Journal of Population Economics* 7(4): 351–378. doi: 10.1007/bf00161472.

Collins, Randall. 2010. "A Dynamic Theory of Battle Victory and Defeat." *Cliodynamics*, 1: 3–25.

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

Goldstone, Jack A. 1991. *Revolution and rebellion in the early modern world*. Berkeley: University of California Press.

Ibn, Khaldun. 1958. *The Muqaddimah: An introduction to history*. Translated by Franz Rosenthal. New York: Pantheon Books.

Krumhardt, Kristen M. 2010. "ARVE Technical Report #3: Methodology for world-wide population estimates: 1000 BC to 1850." École Polytechnique Fédérale de Lausanne, Dept. of Environmental Engineering, ARVE Research Group. http://arve.epfl.ch/technical_reports/ARVE_tech_report3_pop_methods.pdf.

McEvedy, C. and R. Jones, 1978. *Atlas of World Population History*. London: Penguin Books Ltd.

McNeill, William H. 1984. "Human Migration in Historical Perspective." *Population and Development Review* 10(1): 1–18. doi: 10.2307/1973159.

Nefedov, Sergey A. 2013. "Modeling Malthusian Dynamics in Pre-Industrial Societies. " *Cliodynamics* 4: 229–240.

Turchin, Peter. 2003. *Historical Dynamics: Why States Rise and Fall*. Princeton University Press.

Turchin, Peter and Thomas D. Hall. 2003. "Spatial synchrony among and within world-systems: insights from theoretical ecology." *Journal of World Systems Research* 9(1): 37–64. doi: 10.5195/jwsr.2003.248.

Turchin, Peter. 2005. "Dynamical Feedbacks between Population Growth and Sociopolitical Instability in Agrarian States." *Structure and Dynamics* 1(1): Article 3.

Turchin, Peter, Jonathan M. Adams, and Thomas D. Hall. 2006a. "East-West Orientation of Historical Empires and Modern States." *Journal of World-Systems Research* 12(2): 219–229. doi: 10.5195/jwsr.2006.369.

Turchin, Peter and A. Korotayev. 2006b. "Population Dynamics and Internal Warfare: a Reconsideration." *Social Evolution and History* 5(2): 121–158.

Turchin, Peter. 2006c. *War and Peace and War: The Life Cycles of Imperial Nations*. New York: Pi Press.

Turchin, Peter. 2009. "A theory for formation of large states." *Journal of Global History* 4(2): 191–217. doi: 10.1017/s174002280900312x.

Turchin, Peter and Sergey Nefedov. 2009. *Secular Cycles.* Princeton University Press. doi: 10.1515/9781400830688.

Turchin, Peter and Sergei Gavrilets. 2009. "Evolution of Complex Hierarchical Societies." *Social Evolution and History* 8(2): 167–198.

Turchin, Peter, Thomas E. Currie, Edward A. L. Turner, and Sergei Gavrilets. 2013. "War, Space, and the Evolution of Old World Complex Societies." *PNAS* 110(41): 16384–16389. doi:10.1073/pnas.130882511.