

Generating State-Size Distributions: A Geopolitical Model*

Lars-Erik Cederman
Swiss Federal Institute of Technology (ETH)
Center for Comparative and International Studies (CIS)
Seilergraben 49
8092 Zurich Switzerland
+41-1-6326759
lcederman@ethz.ch

Version 1.01: Comments welcome!
September 18, 2003

Abstract:

While International Relations scholars have ignored the issue of state size, economists explain the territorial extent of states as an optimal outcome given various constraints in analogy with the theory of the firm. Focusing on full distributions rather than average sizes, this paper adopts a systemic, generative perspective supported by agent-based modeling. It is concluded that empirical state sizes are log-normally distributed. Given this fact, I attempt to reconstruct such size distributions relying on a geopolitical model. This reconstruction task becomes possible by adding a mountainous terrain that imposes a variety of logistical obstacles to conquest processes.

*) Paper prepared for presentation at the Agent 2003 conference on “Challenges in Social Simulation”, University of Chicago, October 2-4, 2003, and at University of California, San Diego, January 12, 2004. I would like to thank David Lake and his colleagues for the permission to use their new dataset.

Introduction

Polarity has traditionally played a central role in International Relations (IR) theory. Scholars have engaged in lengthy and inconclusive debates about whether bipolarity or multipolarity is more likely to produce geopolitical stability. Yet, despite the attention paid to the number of states, little has been said about their territorial extension. This is surprising, because states are the key actors of most theories, and size is perhaps the most obvious attribute of any organization.

By contrast, economists have spent much more time accounting for state size. Drawing on the classical theory of the firm, their individualistic approach often downplays geopolitics, however. Moreover, it focuses on equilibrium outcomes at the expense of historical developments. While fitting neatly into a textbook microeconomic perspective, such scholarship is almost totally disconnected from traditional IR theory,

In this paper, I address the question of state size square on without losing sight of geopolitics. I do so by adopting a systemic, distributional perspective that is supported by computational modeling. Size distributions contain much more information about a system than polarity, which is nothing more than a single number. Most importantly, distributions may be interpreted as “footprints” of the underlying mechanisms that generate them, thus helping us to explain not only the states’ size but their genesis and behavior in more general terms.

As a starting point, I use a new territorial dataset to show that real-world state sizes are log-normally distributed. Because macro-historical mechanisms cannot be easily manipulated in quasi-experimental terms, I then introduce a computational model that I developed for other purposes to see if it is possible to reconstruct the empirical distributions within that framework (Cederman 2003). It turns out that the original model performs quite poorly in this respect. Adding a map of geographic obstacles, however, takes us much closer to the historical benchmark. This finding quantifies what geopolitical theorists have long suspected, i.e. that rugged terrain and communication technologies have a direct impact on the scale of political organization.

Existing theories of state size

Theorists in IR have spilt a lot of ink arguing about the consequences of specific polarity configurations. Primarily, the controversy concerns the impact of different polarity structures on the level of conflict. Whereas classical realists, such as Morgenthau (1985), maintain that multipolarity increases systemic stability, Waltz’s (1979) influential neorealist theory contends that bipolar systems should on average be more stable. Yet, another group of scholars insists that it is unipolarity that reduces the level of conflict most effectively (e.g. Organski 1968; Gilpin 1981).

Unfortunately, partly due to terminological confusion regarding whether polarity refers to the number of great powers or the number of great-power alliances, several decades’ research has failed to produce conclusive evidence in one or the other direction (Geller and Singer 1998). At a deeper level, the lack of progress probably relates to the fact that the international system, like other complex systems, does not lend itself to *ceteris paribus* reasoning. Quite on the contrary, without endogenizing polarity, we stand little hope of uncovering the net effects of different systemic structures (if there are any). The problem with the traditional research is that it treats the number of states (or clusters of states) as exogenous while studying the effects thereof.

In my previous research, I have attempted to overcome this theoretical difficulty by generating systems with different polarity as endogenous outcomes (see Cederman 1997).

In doing so, I discovered that, contrary to prevalent expectations, defensive systems based on technology or alliances have a tendency to degenerate into unipolarity. While this research represents a step forward in terms of endogeneity, it fails to address the issue of state size directly. As I have argued above, IR theory needs to investigate this fundamental property, especially because it is easily observable. To this date, there are very few attempts to address the issue within this field, though Russett (1968) constitutes an early exception to which we will have reason to return below.

To find research that attempts to account for the differing sizes of states in explicit terms, it is necessary to turn to economics. The classical theory of the firm provides an obvious analogy on which political economists base their theoretical efforts. In a pioneering article, Bean (1973), postulates the existence of a tradeoff between economies of scale and various constraints. In abstract terms, the economies of scale relate to the provision of public goods, such as order and security, health, transport and communications, as well as the means to finance these goods through tax collection. Yet, at larger distances from the political center, increasing costs start mounting, which ultimately turns economies of scale into decreasing returns.

There are several plausible sources of such constraints, including geographical obstacles as well as “cultural distance” that render the task of governance difficult due to lacking inter-ethnic loyalty. Bean (1973) further assumes that military and logistical technology determines the location of the tradeoff in given historical situations. Over time, technological progress is prone to expand the reach of effective rule. For example, the emergence of the territorial state in early modern Europe can be seen as an outgrowth of advances in military technology (e.g. Tilly 1985; Downing 1992).

More recently, however, economists have tended to play down such geopolitical factors, without abandoning the underlying structure of the tradeoff (Wittman 1991). In an oft cited article, Friedman (1977) suggests that territorial state sizes reflect an optimal allocation of tax collection based on land and labor. Further privileging economic over security-related determinants, Alesina and Spolaore (1997; forthcoming) assume culture to be uniformly distributed in space, and postulate “preference heterogeneity” resulting from cultural differences. Inspired by the events following the end of the Cold War, including the collapse of multi-ethnic communist states, their models suggest that state size depends crucially on regime type. Whereas democratic states are expected to be too small compared to the optimal size of public goods provision, the opposite holds for non-democratic states. Focusing on the impact of federalism, Hiscox and Lake (2001) offer an interesting elaboration of this theme. Moreover, Alesina and Spolaore (forthcoming) argue that openness to trade tends to reduce state size as well.

Whereas these political-economy studies differ as regards the weight put on specific factors, they share a staunch commitment to an equilibrium perspective. This is hardly surprising given the importance of microeconomic theory of the firm as the source analogy. Yet, the attempt to derive the scale of state-formation as an optimal outcome forces these analysts to make a number of tacit assumptions that limit the search for mechanisms that determine state size.

First, their perspective is fundamentally static. The organizational configuration of states is expected to adjust quickly to the prevailing conditions. In other words, political economists typically assume history to be “efficient” rather than path-dependent (March and Olsen 1984). Such an analysis focuses more on cross-sectional comparisons while ignoring the mechanisms that generated the outcomes in the first place.

Second, the equilibrium formalism of microeconomics puts the spotlight primarily on the individual level. Almost without exception, economists attempt to derive explanations for state size through comparative statics analysis that isolates the country cases before comparing them. But given the interconnected nature of the international system, and the zero-sum character of territory, this explanatory strategy has obvious limitations.

Third, the classic theory of the firm is based on rational choices, thus drastically overstating the voluntaristic aspects of macro-level processes. While states are ultimately created through human choices, it does not follow that they result from long-term plans reflecting optimal choice (cf. Tilly 1985; Spruyt 1994). Furthermore, economic rationales are typically privileged, including implausible, complex calculus involving maximization of vote shares and tax revenue. While far from all territorial change results from war and conquest, the economic analysis usually sweeps geopolitics under the rug.

Fourth, and finally, the search for closed-form equilibrium solutions usually forces the analysts to impose the wildly inaccurate assumption of uniform sizes. Though some attempts to relax this assumption have been made (see e. g. Alesina and Spolaore 1997), the microeconomic literature on state size has usually nothing to say about the entire distribution.

It seems appropriate to make a radical break from these assumptions. In the following, I will propose an approach that is self-consciously dynamic, systemic, structural, and distributional. While it has never been followed up in IR scholarship, Russett (1968) sketched exactly such a research program. Building on a dataset on population size, he discovered that his population data conforms with log-normal distributions. Given this insight, “it is possible from the descriptive material on the size distribution of nations to make some inferences about how the distribution got that way” (p. 306). Indeed, while providing a formal model and shifting the focus from population to territorial sizes, Russett’s approach resembles strongly the generative logic that I follow in this article.

Empirical state-size distributions

The first challenge confronting a distributional perspective is to explore empirical data in order to determine what type of distribution is operating. Fortunately, there is a rich literature on size distributions in various domains. For our purposes, it is particularly interesting that some economists, deviating from the microeconomic orthodoxy, have studied the systemic statistical patterns relating to firm sizes. The locus classicus of this scholarship can be traced back to “Gibrat’s Law,” or “the law of proportionate effect,” which states that multiplicative random walks tend to generate log-normal distributions (Sutton 1997). More simply put, this applies to processes that entail that an organization’s growth is proportionate to its size. Formally, log-normality implies that size S is distributed according to the following principle:

$$\log S \sim N(\mathbf{m}, \mathbf{s})$$

were \mathbf{m} is the mean and \mathbf{s} the standard deviation (see Aitchison and Brown 1957; Crow and Shimizu 1988).

Subsequent empirical research has shown that there are other skew distributions that perform well as descriptive statistics of firm sizes. Following Simon and Bonini (1958), power laws are often mentioned as plausible candidates. Such distributions have a “thicker tail” that reflects a higher frequency of very large firm sizes (Axtell 2001). Wars,

measured in terms of casualty numbers, are power-law distributed (Richardson 1960; Cederman 2003).¹

How can we tell log-normal distributions apart from power laws? The easiest way is to plot the logged cumulative frequency (c.d.f.) against logged sizes. Whereas power laws should appear as straight lines in such a frequency diagrams, log-normal distributions taper off for large sizes, and therefore exhibit significant bending.

Luckily, I have been able to use a new dataset on the territorial size of states generously provided by Lake and O'Mahony (2002) and Hiscox and Lake (2001). Based on information drawn from several data sources, including COW and Polity III, their data base covers the period between 1815 and 1998 excluding the world wars and the colonial empires.

Figure 1 and 2 display the frequency diagrams for the first and the last year of the sample. The diagrams depict the logged, converted c.d.f., $\log \Pr(S > s)$, as a function of logged size $\log s$. Throughout this paper, I use logarithms with a base of ten. All empirical state sizes are measured in terms of square kilometers. Intuitively speaking, the function indicates the probability that there will be a size of an even larger size. For very small sizes, this probability is close to one, but as the size increases, it falls quickly to very small values.

[Figure 1: Empirical state sizes in 1815]

[Figure 2: Empirical state sizes in 1998]

Visual inspection immediately reveals that the empirical state sizes do *not* follow a power law due to the obvious curvature of the data. To see if they are log-normally distributed instead, I generated such curves using maximum-likelihood estimation, as indicated by the Figures 1 and 2 (see Aitchison and Brown 1957). This numerical analysis shows that for the two years shown, a log-normal distribution captures the observations quite well.

Already in 1815, the data conforms quite accurately to a log-normal distribution. For this year, the estimation based on 34 observations yields

$$\log S \sim N(4.98, 1.02)$$

with a mean absolute error (MAE) of 0.048.² In 1998, the fit is even more impressive, as reflected by the lower MAE of 0.028 based on as many as 154 state sizes. Figure 2 demonstrates that the empirical curve approximates log-normality very well. At this point, the estimated curve approximates the following distribution:

$$\log S \sim N(5.31, 0.79).$$

In addition, I estimated log-normal distributions for each year in the sample in order to assess the log-linear fit for the entire sample period. Figure 3 tells us that the MAE values remain relatively stable around 0.04 throughout the 19th century. In the interwar period,

¹ Of course, there are other skew distributions, such as the Yule distribution (see Ijiri and Simon 1977; Reed n. d.).

² The median absolute error offers an intuitive measure of the fit of the curves in the units of the observations. Note that the errors are calculated based on the logged axes. Other standard measures could have been used as well, such as the mean square error. Yet, the most important criterion is the visual shape of the point cloud. Small MSE values may hide systematic deviations from log-normality.

the fit worsens somewhat, but after the early post-WWII period, a stable regime sets in with even lower MAEs around 0.03. With the exception of the years immediately succeeding WWII, which are characterized by acute instability, the data seem robustly log-normally distributed through the entire sample.³

[Figure 3: Mean absolute errors (MAEs) for log-normal distributions, 1815-1998]

As hinted by our two snapshots, the m values of these distributions have also fluctuated somewhat over time (see Figure 4). Until about 1880, the trend pointed toward a quick increase of state size. After this point, however, the growth process turned into a slow reduction from top values of almost 5.6 down to 5.31 for the last sample year. In fact, since the last 19th century, there has been an almost linear decrease of the distributional mean (see also Lake and O'Mahony 2002).

[Figure 4: Estimated m values for log-normal distributions, 1815-1998]

Unsurprisingly, the distributional shifts in state size are intimately linked to polarity. Figure 4 displays the number of sovereign states according to the data set. Here the development is quite stable during the 19th century. After a slight increase from the initial 34 states in 1815, the numbers fluctuate around 50 for the remaining century. After 1900, however, the number of states grows explosively in the course of the 20th century up to 154 in 1998.

[Figure 5: Polarity of the international system, 1815-1998]

At this point, it is only possible to speculate about the precise reasons for this explosive increase in polarity. In general terms, however, it seems reasonable to postulate that the advent of mass politics as expressed by nationalism and democracy made the key difference. Together with principles such as popular sovereignty and self-determination, these processes profoundly changed the nature of boundaries in the international system provoking decolonization and other cases of imperial decay.

Despite these changes over time, it appears that Russett's insights about population size can be readily generalized to territorial data. With very few exceptions, the historical pattern confirms that territorial state sizes are indeed log-normally distributed with a very small margin of error.

What explains the strikingly good fit? What theoretical inferences can be drawn from it? The obvious problem is that there is potentially an infinite number of mechanisms that are capable of producing log-normal laws (see Russett 1968, p. 315). Yet, these empirical findings are useful because they can be used as an explanatory target. After all, far from every model is capable of generating the patterns in question. What we can say then is that any systemic theory of state size worth its salt, or possibly even any general macro-theory of International Relations, has to reproduce this distributional footprint in order to claim quantitative accuracy.

Nevertheless, in enormously complex settings, such as the Westphalian state system it is hard to match mechanisms with aggregate outcomes. For obvious reasons, counterfactual substitutions become increasingly difficult as soon as we distance ourselves from the

³ If the observations are limited to Europe only (excluding Russia and the USSR), small sample size means that the MAE values fluctuate more than in the entire international system. While starting just above 0.06, there are deviations in the late 19th century that transgress the 0.1 limit. By 1998, however, the MAE value went down to 0.055.

historical path. Indeed, it is implausible that as fundamental properties as state sizes flow from superficial, short-term processes.

Computational modeling provides tools that promise to alleviate this dilemma. If we cannot conduct experiments in world history, at least we can recreate simplified, artificial worlds which lend themselves to being experimentally manipulated (Cederman 1997, Chap. 3). Within in such framework, it becomes possible to investigate if there are specific mechanisms or conditions that generate the empirically observed patterns.

Geosim: An agent-based model

During the past decade, I have developed a computer-based, geopolitical laboratory for various analytical purposes (Cederman 1994; 1997). Developed in the agent-based tradition of Bremer and Mihalka (1977), these models share a common architecture featuring raw Hobbesian power competition among perfectly sovereign states. The current study uses a version, called GeoSim, that was developed to regenerate power-law distributed wars (Cederman 2003). As opposed to the earlier versions which were coded in Pascal, GeoSim was implemented using the Java-based toolkit Repast.

Does GeoSim produce log-normal state sizes, or is some addition needed to achieve this goal? Before attempting to answer this question, however, it is necessary to introduce GeoSim's main principles. Because Cederman (2003) describes that model in great detail, I will limit the current exposition to a summary of the previous version, and then focus on the incremental changes that were added in this paper.

The standard initial configuration consists of a 50 x 50 square lattice populated by about 200 composite, state-like agents interacting locally. Because of the boundary-transforming influence of conquest, the interactions among states take place in a dynamic network rather than directly in the lattice. In each time period, the actors allocate resources to each of their fronts and then choose whether or not to fight with their territorial neighbors (see Figure 6). In the grid, the lines correspond to state borders, and the dots or rings, to the capitals.

[Figure 6: A multipolar sample system at time 5000]

All states use the same “grim-trigger” strategy in their relations. Normally, they reciprocate their neighbors' actions. Should one of the adjacent actors attack them, they respond in kind without relenting until the battle has been won by either side or ends with a draw. Unprovoked attacks can happen as soon as a state finds itself in a sufficiently superior situation vis-à-vis a neighbor. Set at a ratio of three-to-one with respect to the locally allocated resources, a stochastic threshold defines the offense-defense balance.

Due to the difficulties of planning an attack, actors challenge the status quo with a low probability. When the local capability balance tips decisively in favor of the stronger party, conquest results, implying that the victor absorbs the targeted unit. This is how hierarchical actors form. If the target was already a part of another multi-province state, the latter loses its province. Successful campaigns against the capital of corporate actors lead to their complete collapse.

Territorial expansion has important consequences for the states' overall resource levels. After conquest, the capitals of conquered territories are able to “tax” the incorporated provinces including the capital province. As shown in Figure 7, the extraction rate depends on the loss-of-strength gradient that approaches one for the capital province but that falls quickly as the distance from the center increases (Boulding 1963; Gilpin 1981,

p. 115). This function also governs power projection for deterrence and combat. Given this formalization of logistical constraints, technological change is modeled by shifting the threshold to the right, a process that allows the capital to extract more resources and project them power farther away from the center. In the simulating runs reported in this paper, the transformation follows a linear process in time. (Note that in the grids, states that have undergone technological change at least once have their capitals depicted as rings rather than as dots, see Figure 6.)

[Figure 7: Technological change as shifts of the loss-of-strength gradient]

Together all these rules entail four things: First, the number of states will decrease as the power-seeking states absorb their victims. Second, as a consequence of conquest, the surviving actors increase in territorial size. Third, decentralized competition creates emergent boundaries around the composite actors. Fourth, once both sides of a border reach a point at which no one is ready to launch an attack, a local equilibrium materializes.

State-size distributions in the unmodified GeoSim model

How could the agent-based model shed light on the historical record of state sizes? It may seem that the simulation model has little to do with the empirically observed process described above. Whereas the model produces a steady fall in polarity, the dataset of Lake and his colleagues points in precisely the opposite direction.

If we expand the time perspective somewhat, however, it becomes clear that within the Westphalian state system, polarity has actually decreased dramatically over the centuries. Though there are no precise figures, historians have counted about a thousand independent political units in the Middle Ages (Jones 1981, p. 106). In early modern Europe, there were still half a thousand such units (Tilly 1975). What followed was a phenomenal consolidation down to about 25 states in the 19th century (i.e. half of the polarity reported for that period above).

Rather than trying to model the last two hundred years of history, I limit the goal of this paper to reproduce state sizes as they appeared at the beginning of the empirical sample period, i.e. during the first half of the 19th century, before nationalism and democracy started to have geopolitical repercussions. This seems to be a reasonable limitation, because GeoSim in its current form does not attempt to trace modern, participatory politics, be it in terms of nationalism or democracy.⁴

Therefore, the high initial polarity characterizing the simulation runs makes sense. The main target, then, is to generate state-size distributions that are log normal at a level of accuracy comparable with the estimated curves reported in the empirical section. To make a comparison possible, it is important to produce state systems with roughly the same number of states as in the 19th century, i.e. about 50.

The experimental procedure goes as follows:

- (1) Run a batch of 15 runs for 10,000 periods.⁵
- (2) Calculate the mean polarity for each time period of all the runs.⁶

⁴ For attempts to extend GeoSim to such settings, see Cederman (2001; 2002).

⁵ Strictly speaking, all runs include an initial period lasting 500 periods before measurement starts. Thus, the total run time amounts to 10,500 periods. All run times are indicated including the initial period.

- (3) Select a time point t^* with an average of approximately 50 states.⁷
- (4) Estimate log-normal distributions for all the 15 runs at t^*
- (5) Select a representative run with the median mean absolute error.
- (6) Use this value, and visual inspection, to evaluate the log-normal fit.

Applied to the 15 runs of the standard model in Cederman (2003), Step 3 tells us to stop at time period 5000. At this point, the representative run with the median fit corresponds to the 55-state system shown in Figure 6 above.

What does the synthetic state-size distribution look like? A glance at Figure 8 suggests that it does not resemble a log-normal pattern at all. Rather than coinciding with the estimated c.d.f., the point cloud intersects the curve, indicating that large states are over-represented and that there is a scarcity of intermediate size states. The almost vertical drop for large state sizes suggests that these units are too similar to conform with a log-normal distribution. The relatively high mean absolute error at 0.086 confirms the deviation from the empirical target. Compared to the distributions estimated in the 19th century, this value is at least twice as high as in most of those cases.

[Figure 8: Representative size distribution at time 5000]

Careful scrutiny of the separate reveals a recurrent pattern featuring historical breaking points that allow a small number of “great powers” to out-grow the rest of the states. But this happens more or less simultaneously, thus creating historically unrealistic equality in terms of the territorial sizes. This pattern is typical for most of the runs. The histogram in Figure 9 shows that the run with median fit is indeed representative of the whole ensemble of 15 simulations. Some error values more than triple the empirically observed rates.

[Figure 9: Distribution of deviations from log-normality, MAE]

In sum, there can be no doubt that GeoSim, in its standard configuration, fails to generate realistically distributed state sizes. The obvious question arises, whether it would be possible to modify the system such as to bring the model’s output in closer harmony with the empirical patterns?

Adding rugged terrain

The computational experiments presented in the previous section tell us that a higher degree of geopolitical diversity is needed. Virtually all theories of state size postulate the presence of constraints that impose increased costs as state sizes grow larger, although the identity of these obstacles vary according to the theory. In keeping with the geopolitical focus on the GeoSim framework, I let mountainous and inaccessible terrain slow down conquest and territorial expansion. Together with cultural differences, this is the factor that is mentioned most often in explanations of state size (see e.g. Bean 1973). In his comprehensive study of European state formation, E. L. Jones (1981, p. 106) presents the basic logic:

Belts of difficult terrain lying between the core-areas, and ancient ethnic and linguistic apartheid dating from early folk movements and settlement history,

⁶ To speed up the simulations, state-size distributions are computed only every 500 time steps.

⁷ I do not count “atomic” states comprising only one unit, because such entities reflect a lower “time resolution” in that they occur as a part of state break-ups. See the rules of “structural change” described in Cederman (2003).

helped to maintain the individuality of political units. Amalgamation went so far but no farther; never to a single empire. Amalgamation costs were high. Major natural barriers protect several parcels of territory the size of modern nation-states and the more durable polities to fit the framework and there stop.

While the GeoSim framework can be adopted to feature cultural differences and nationalism (see Cederman 2002), this paper focuses exclusively on geographic obstacles. On the whole, Jones' theoretical perspective dovetails nicely with the modeling project presented here.

Optimal-size solutions for European states cannot be worked out as simple geometry. The spaces on the board have different values like those in the game of Monopoly and capturing and amalgamating some of them is exceedingly expensive. Where terrain did not provide much protection, units tended to disappear in takeovers by their neighbours (p. 107).

Rather than trying to account for specific state sizes, Jones contents himself with offering a “lower-bound theory of European state formation in which other forces decide the precise outcome, but the selection of the nucleus of the rising state will be from among the richer potential cores” (p. 109). Again, this corresponds very closely to the explanatory ambition of the current paper.

The task, then, is to modify GeoSim in order to mimic real-world geographical constraints. This can readily be done by creating an artificial topological map that allocates a “height” to each cell in the grid. First, a tunable number of mountain peaks is distributed randomly across the grid. Then a diffusion process connects these peaks with their surrounding sites, thus creating relatively smooth mountainous terrain gently sloping down to the plains.⁸

Having described the initialization of the landscape, we need to consider the behavioral implications of the rugged terrain. The key to these modifications resides in the notion of effective, rather than geographic, distance. Whereas the base model uses Cartesian distances as input to the logistical curve presented in Figure 7 above, the modified version replaces this measure with one that takes the difficulty of the local terrain into account. The additional obstacle added to the distance per unit depends on the “altitude” of the mountains where the peaks represent the maximum logistical penalty. In the simulation below, we will assume initially this parameter to be three. Each time a province is conquered, the effective distance from the capital is calculated based on the effective distance of the conquering province, adding the terrain-corrected value of the conquered area. If a mountain peak is conquered, this means that three effective distance units rather than one is added to the accumulated distance from the capital. Apart from this change, the logistical distance curve shown in Figure 7 remains unchanged. It should

⁸ To be more precise, the terrain module allocates heights to every cell in the 50 x 50 grid. The algorithm starts by creating a random selection of mountain summits constituting a fraction `propMountains = 0.05` of all unitary cells. Then a recursive algorithm is run repeatedly `timesSmooth = 20` times smoothening the height of the surrounding cells. This “brush” continues to a cell in the von Neumann neighborhood with a probability of `prVisit = 0.8`. For each “visit”, the algorithm sets the neighboring cell to a weighted average (`propSmooth = 0.7`) of the initial cell and its previous value. From an intuitive standpoint, this algorithm is similar to water flowing from the mountain peaks dragging with it soil that leads to a smoother landscape around the summits.

be stressed again that capitals have to cope with these constraints both when allocating and projecting resources.⁹

As an illustration, Figure 10 displays a grid with a geographical setting of this type. The gray shades correspond to the altitude of the virtual mountains. This particular snapshot describes the situation in time period 7000. At this point, state formation has already generated a number of larger, compound states. As expected by Jones (1981), the largest states are located in the more easily accessible basins (which are depicted as the brighter areas) whereas the smaller units can be found in the mountains and the periphery of the system. Moreover, as a rule, state borders tend to coincide with the mountain ranges.

[Figure 10: Snapshot from system with rugged terrain at time 7000]

What difference do geographical obstacles make in terms of state-size distributions? Following the same experimental steps as in the previous section, I concluded that the 15 replications generated about 50 states at time period 7000. It is not surprising that t^* is somewhat higher for the mountainous terrain than without, because the additional obstacles can be expected to slow down state formation somewhat. Again, the run associated with the MAE was selected and the distribution estimated. Figure 11 presents the results of this procedure.

[Figure 11: Representative size distribution for system with terrain at time 7000]

It is immediately apparent that this distribution comes much closer to the empirical benchmark. While there are some deviations for intermediate-size states, on the whole, the observations conform roughly with the estimated curve. This is reflected in a much lower mean MAE of 0.050 than the one obtained with the base model. The histogram shown in Figure 12 helps us to gauge the fit in all the fifteen runs. A comparison with Figure 9 reveals that the fit for the geography-dependent runs approximate log-normality much better than those of the base model.

[Figure 12: Distribution of log-normal fits in system with terrain at time 7000]

The experimental procedure outlined above prescribes the selection of a particular polarity level. While the choice of about 50 states reflects a specific historical benchmark, the finding that terrain helps generate empirically realistic state sizes would be potentially fragile if it hinged on the number of states in the system. Therefore, I compiled an additional graph which plots the average of the mean absolute errors in all the fifteen runs over time (see Figure 13). This allows a dynamic comparison of the two experimental configurations, with and without terrain.

[Figure 13: Mean value of log-linear fits over time with and without terrain]

Based on the findings shown in Figure 13, it must be concluded that the runs with terrain (marked with dots) produce a persistently tighter fit than without it, in fact almost at the level of real-world data. Moreover, on average, the discrepancies remain steady throughout the course of the simulations. In the base model, however, the runs shift decisively away from the log-normal curve between 3000 and 4000 time steps into the simulations. This effect is precisely what was captured in the snapshot in Figure 8.

⁹ This algorithm calculates distances based on the square grid, as if the system were a big Manhattan, rather than as the crow flies.

Yet, it is obviously premature to draw any far-reaching conclusions about the influence of terrain on state size. All that I have done so far is to show that in a particular, artificial world, the finding can be upheld. Because most of the parameters have been intuitively tuned in order for the model to behave realistically, there is still an urgent need for sensitivity testing and further empirical calibration.

Preliminary sensitivity analysis

Due to space reasons, this paper can offer only very limited insights into the robustness of the computational findings. This section explores three theoretical dimensions. We start by considering different levels of terrain obstacles. Then the attention shifts to the process of technological change. Finally, on a preliminary basis, I attempt to reconstruct a size distribution that approximates the current state of the international system.

To facilitate analytical comparisons, Table 1 provides an overview of these experiments. Corresponding to a set of fifteen runs, each line contains information about representative distributions selected according to the six-step procedure, and statistics about the whole set of runs. The shaded lines refer to the two configurations that have already been discussed. Whereas Line 1 is associated with the base model, Line 3 introduces rugged terrain in the standard configuration with an altitude of three.

[Table 1: Results from sensitivity analysis]

Focusing on the difficulty of the terrain, together Lines 2, 3, and 4 tell us that the other levels of obstacles produce similar, if not equally impressive results. Yet, in all cases, the log-normal fit clearly improves over the base runs in Line 1. Thus, other things being equal, the main finding appears to hold reasonably robustly.

The second dimension highlights technological change. Tuned to provide power laws over a large size range in Cederman (2003), the standard rate shifts the loss-of-strength gradient by as much as 20 units from an initial two units during the course of each simulation (see Figure 7). To explore a less dramatic geopolitical transformation, the next group of runs features half of that rate, i.e. a shift by 10 units. While the runs without terrain, shown in Line 5, differ very little from the base runs in Line 1, geographical constraints contribute even more strongly to generate realistic output. Lines 6, 7, and 8 reveal that for each level of ruggedness, the log-normal fit improves compared to the first set of runs with a high rate of technological change. In fact, for mountain heights of three, the MAE value 0.043 of the representative run is half of the one without terrain.

Still these results fall short of reproducing state sizes comparable to those in today's international system. Though this is not the prime scenario due to the model's limited suitability in such scenarios, it is still interesting to see how far the model can be pushed. Because the real-world distribution in 1998 has as many as 154 states, it seems reasonable to generate around 150 rather than 50 states at time t^* . Yet, in order to produce polarity levels of this magnitude, it is necessary to increase the dimensions of the grid from 50 x 50 to 75 x 75. Lines 9 and 10 report on the results. Here terrain makes an even marked difference than in the smaller grids. With a median MAE of 0.032, the topologically modified runs come very close to the empirical fit (which has a MAE of 0.027). Yet, despite the low MAE value, the shape does not live up to log-linearity because the points exhibit a slightly exaggerated curvature, especially for large states. Whether this depend on the absence of nationalism or a fragility in the present model requires further research that goes beyond the scope of this study.

Before considering such extensions, however, it is very important to check whether the power-law result in terms of war sizes of Cederman (2003) remains robust despite the topological alterations introduced in this paper. It should be recalled that Line 1 produces exactly the same war behavior as the earlier paper. A quick check of the standard model in Line 2 suggests that scaling survives the changes without problems. Inspection of all the 15 log-log plots (not shown here) confirms that linearity is upheld across the board. Numerically the result is indistinguishable from that reported in Cederman (2003), because the median R^2 is as high as 0.994 as opposed to 0.991 for Line 1, and the range of R^2 values has a higher minimum: [0.983 0.997]. As might be expected, the slopes become steeper (a median of -0.69 rather than -0.55).¹⁰ It can thus be concluded that the geographic modification generate realistic sizes of both wars and states.

Having investigated the findings' robustness along a few selected dimensions, we are still not in a position to draw any firm conclusions about the independent effect of logistical constraints on state sizes. Further sensitivity analysis and empirical calibration are clearly needed. A promising avenue of analysis would build on GIS tools and other sources of geographic data to derive empirical measures of terrain and communications. Such data could then be used for calibration of the model (cf. Lake and O'Mahony 2002). For example, investigations of the entire international system appear seriously incomplete without more attention paid to naval warfare and sea communications (e.g. Rasler and Thompson 1989).

Above, I have already conjectured that the decrease in average state size throughout the 20th century, together with the explosive proliferation of independent states, have a lot to do with nationalism and participatory politics: "These nineteenth-century developments differed fundamentally from other histories of territorial consolidation in Europe" (Rokkan 1999, p. 263). To capture such effects, the model needs to be explicitly extended, something that is fortunately within the realm of possibilities. Such an extension requires both a mechanism for secession and a representation of culture and identity (see Cederman 2002). In addition, it would be desirable to develop empirical measures of the colonial empires' territorial scope.

Beyond this important challenge, the long-term goal would be to join the economists and other scholars who attempt to explain the size of particular types of states. I would submit, however, that the current systemic approach provides a more solid theoretical, because it allows us to reject models that fail to produce the distributional benchmarks. Moreover, as I have argued, it appears very likely that state size is to a large extent an inherently systemic attribute.

Finally, it should be remarked that elaborations of this type represent one possible strategy. Equally interesting, however, is the idea of radically simplifying the model. Such a project might say more about the logic of the underlying mechanisms, the detailed operation of which is very hard to trace in the GeoSim framework (see e.g. Stanley et al. 1996).

Conclusion

On balance, this study demonstrates that realistic, log-normal state-size distributions can be "grown" artificially. To my knowledge, nobody has ever proposed a model that does this. The main finding confirms the intuition of systemic theories that stress topological

¹⁰ For shock levels at 10, the improvement is equally significant (see Line 7). In this case, the median R^2 value is 0.991 with a range [0.980 0.996] and the median slope -0.72.

constrains as a course of geopolitical diversity. Furthermore, the model generates power-law distributed war sizes, thus hitting two important empirical targets at the same time.

Of course, the current paper is not the definitive word on the topic of state size. Quite on the contrary, one would hope that it will inspire others to take on the challenge of modeling macro-historical processes with computational tools. For this approach has many obvious advantages. First, as a contribution to IR theory, it goes well beyond the traditional debate about polarity, with its unclear definitions and failure to endogenize the number of states. As a complement to traditional quantitative studies and rationalistic model-building, computational models of this kind can serve an important purpose in checking the empirical plausibility and internal consistency of systemic theorizing in IR.

Second, agent-based modeling also constitutes a useful alternative to individualist theories of state size, especially to those that regard this property as an equilibrium outcome. By offering an explicit representation of the dynamic mechanisms constituting non-equilibrium processes, the generative approach sheds more light on the sources of the empirical patterns than does comparative statics analysis. Since longitudinal data is available, it would be a shame not to take advantage of this information in theory-building.

Third, by introducing a new way of experimenting with what's profoundly unalterable in the real world, the current computational approach will hopefully revive long neglected geopolitical scholarship and put macro-sociological analysis on a surer footing. Without the "accounting mechanism" of agent-based modeling, it is hard to assure that intuitively compelling arguments be connected with observed macro patterns.

References

- Aitchison, John, and James A. C. Brown. 1957. *The Lognormal Distribution*. Cambridge: Cambridge University Press.
- Alesina, Alberto, and Enrico Spolaore. 1997. On the Number and Size of Nations. *Quarterly Journal of Economics* 112:1027-1056.
- Alesina, Alberto, and Enrico Spolaore. forthcoming. *The Size of Nations*. Cambridge: MIT Press.
- Axtell, Robert L. 2001. Zipf Distribution of U.S. Firm Sizes. *Science* 293:1818-1820.
- Bean, Richard. 1973. War and the Birth of the Nation State. *Journal of Economic History* 33:203-221.
- Boulding, Kenneth E. 1963. *Conflict and Defense*. New York: Harper and Row.
- Bremer, Stuart A., and Michael Mihalka. 1977. Machiavelli in Machina: Or Politics Among Hexagons. In *Problems of World Modeling*, edited by K. W. Deutsch. Boston: Ballinger.
- Cederman, Lars-Erik. 1997. *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*. Princeton: Princeton University Press.
- Cederman, Lars-Erik. 2001. Modeling the Democratic Peace as a Kantian Selection Process. *Journal of Conflict Resolution* 45 (August):470-502.
- Cederman, Lars-Erik. 2002. Endogenizing Geopolitical Boundaries with Agent-Based Modeling. *Proceedings of the National Academy* 99 (suppl. 3):7796-7303.
- Cederman, Lars-Erik. 2003. Modeling the Size of Wars: From Billiard Balls to Sandpiles. *American Political Science Review* 1 (97):135-150.
- Crow, Edwin L., and Kunio Shimizu. 1988. *Lognormal Distributions: Theory and Applications*. New York: Marcel Dekker.

- Downing, Brian M. 1992. *The Military Revolution and Political Change : Origins of Democracy and Autocracy in Early Modern Europe*. Princeton: Princeton University Press.
- Friedman, David. 1977. A Theory of the Size and Shape of Nations. *Journal of Political Economy* 87 (1):59-77.
- Geller, Daniel S., and J. David Singer. 1998. *Nations at War: A Scientific Study of International Conflict*. Cambridge: Cambridge University Press.
- Gilpin, Robert. 1981. *War and Change in World Politics*. Cambridge: Cambridge University Press.
- Hiscox, Michael J. and David A. Lake. 2001. Democracy and the Size of States. Ms. Department of Political Science, UCSD.
- Ijiri, Yuji, and Herbert A. Simon. 1977. *Skew Distributions and the Sizes of Business Firms*. New York: North-Holland.
- Jones, E. L. 1981. *The European Miracle: Environments, Economies and Geopolitics in the History of Europe and Asia*. 2nd ed. Cambridge: Cambridge University Press.
- Lake, David A., and Angela O'Mahony. 2002. The Incredible Shrinking State: Explaining Change in the Territorial Size of Nations. Ms. University of California, San Diego.
- March, James G., and Johan P. Olsen. 1984. The New Institutionalism: Organizational Factors in Political Life. 78 (3):734-749.
- Morgenthau, Hans J. 1985. *Politics Among Nations*. New York: Knopf Publishers.
- Organski, A. F. K. 1968. *World Politics*. 2nd ed. New York: Knopf.
- Rasler, Karen A., and William R. Thompson. 1989. *War and State Making: The Shaping of Global Powers*. Boston: Unwin Hyman.
- Reed, William J. n. d. The Pareto, Zipf and Other Power Laws. Ms. Department of Mathematics and Statistics, University of Victoria, B. C.
- Rokkan, Stein 1999. *State Formation, Nation-Building, and Mass Politics in Europe: The Theory of Stein Rokkan*. Flora, Peter, ed. Oxford: Oxford University Press.
- Russett, Bruce M. 1968. Is There a Long/Run Trend Toward Concentration in the International System? In *Economic Theories of International Politics*, edited by B. M. Russett. Chicago: Markham.
- Simon, Herbert A., and Charles P. Bonini. 1958. The Size Distribution of Business Firms. *American Economic Review* 48 (4):601-617.
- Spruyt, Hendrik. 1994. *The Sovereign State and Its Competitors: An Analysis of Systems Change*. Princeton: Princeton University Press.
- Stanley, M. H. R., et al. 1996. "Scaling Behavior in the Growth of Companies." *Nature* 379: 804-806.
- Sutton, John. 1997. Gibrat's Legacy. *Journal of Economic Literature* 35:40-59.
- Tilly, Charles. 1975. Reflections on the History of European State-Making. In *The Formation of National States in Western Europe*, edited by C. Tilly. Princeton: Princeton University Press.
- Tilly, Charles. 1985. War Making and State Making as Organized Crime. In *Bringing the State Back In*, edited by P. B. Evans, D. Rueschemeyer and T. Skocpol. Cambridge: Cambridge University Press.
- Waltz, Kenneth N. 1979. *Theory of International Politics*. New York: McGraw-Hill.
- Richardson, Lewis F. 1960. *Statistics of Deadly Quarrels*. Chicago: Quadrangle Books.
- Wittman, Donald. 1991. Nations and States: Mergers and Acquisitions; Dissolutions and Divorce. *AEA Papers and Proceedings* 81 (2):126-129.

Tables

Table 1. Results from the sensitivity analysis

#	Parameters			Representative distribution					Output from all 15 runs				
	dim	tech	terr	t^*	#states	m	s	MAE	mean #states	mean m	mean s	min MAE	max MAE
1	50	20	0	5000	57	1.38	0.48	0.085	49	1.54	0.47	0.046	0.126
2	50	20	2	7000	62	1.51	0.35	0.055	50	1.51	0.47	0.039	0.085
3	50	20	3	7000	43	1.47	0.53	0.050	53	1.45	0.48	0.035	0.088
4	50	20	4	7000	60	1.37	0.48	0.059	52	1.42	0.52	0.045	0.089
5	50	10	0	8500	50	1.52	0.36	0.086	52	1.49	0.41	0.058	0.137
6	50	10	2	10500	53	1.48	0.43	0.049	54	1.50	0.43	0.036	0.084
7	50	10	3	10500	51	1.48	0.46	0.043	59	1.45	0.43	0.031	0.062
8	50	10	4	10500	54	1.39	0.52	0.048	54	1.46	0.46	0.030	0.076
9	75	10	0	7000	172	1.34	0.40	0.092	140	1.40	0.42	0.047	0.117
10	75	10	3	9000	146	1.41	0.44	0.032	154	1.38	0.45	0.020	0.040

Legend:

dim = dimension of grid

tech = rate of technological change

terr = max level of terrain obstacles (mountain altitude)

Shaded lines refer to the runs discussed before the sensitivity analysis

Figures

$\log \Pr(S < s)$




Figure 1. Empirical state sizes in 1815.

$\log \Pr (S > s)$




Figure 2. Empirical state sizes in 1998.

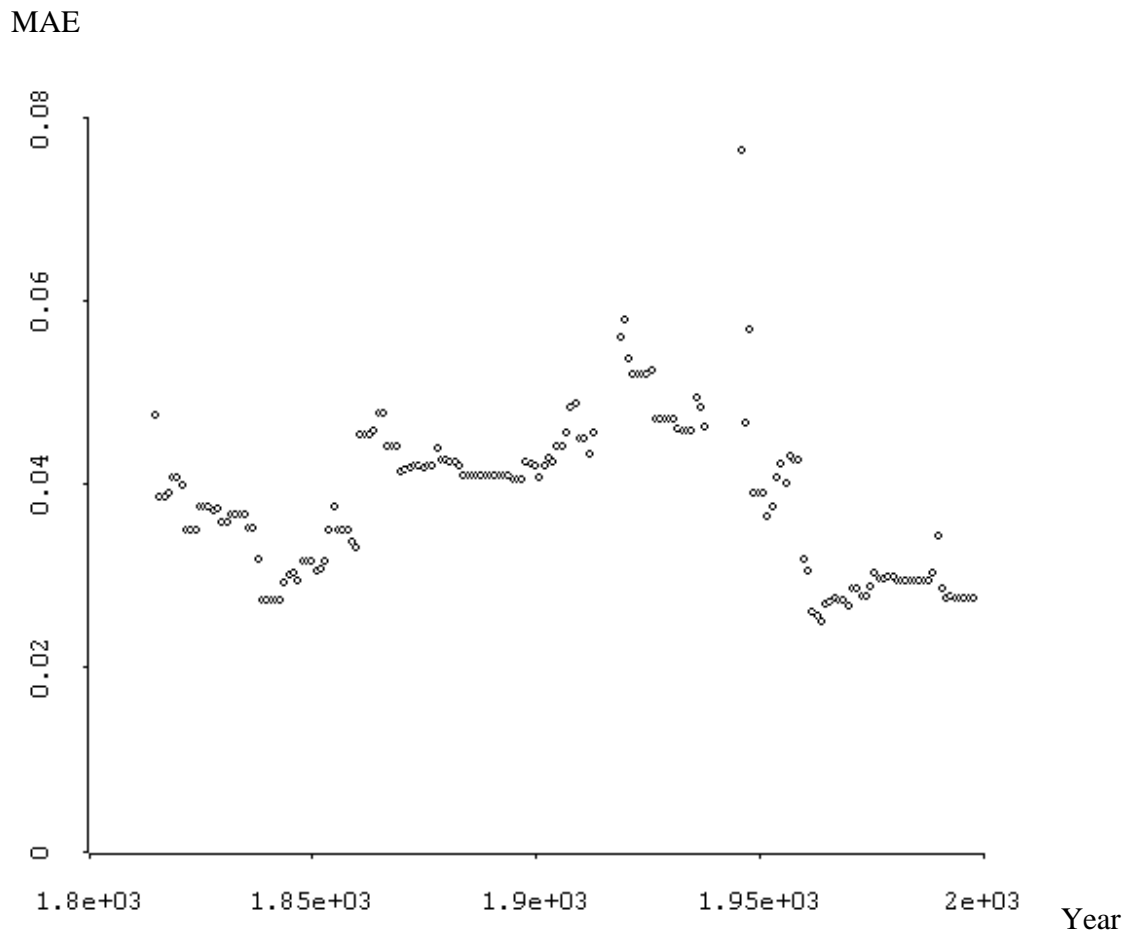


Figure 3. Mean absolute errors (MAEs) for log-normal distributions, 1815-1998




Figure 4. Estimated m values for log-normal distributions, 1815-1998

Polarity




Figure 5. Polarity of the international system, 1815-1998




Figure 6. A multipolar sample system at time 5000

Degree of resource
extraction and projection




Figure 7. Technological change modeled as a shift of loss-of-strength gradients

$\log \Pr(S > s)$




Figure 8. Representative size distribution at time 5000




Figure 9. Distribution of deviations from log-normality, MAE




Figure 10. Snapshot from system with difficult terrain at time 7000

$\Pr(S > s)$




Figure 11. Representative size distribution for system with terrain at time 7000





Figure 12. Distribution of log-normal fits in system with terrain at time 7000.

Mean MAE
in 15 runs



Legend:
Straight line = Systems without terrain
Dotted line = Systems with terrain (height three)

Figure 13. Mean value of log-linear fits over time with and without terrain.