

Research Article

Modelling Hegemonic Power Transition in Cyberspace

Dmitry Brizhinev, Nathan Ryan , and Roger Bradbury 

National Security College, The Australian National University, Canberra, ACT 0200, Australia

Correspondence should be addressed to Roger Bradbury; roger.bradbury@anu.edu.au

Received 14 September 2017; Revised 10 February 2018; Accepted 21 February 2018; Published 11 April 2018

Academic Editor: Carlos Gershenson

Copyright © 2018 Dmitry Brizhinev et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Cyberspace is the newest domain of conflict and cooperation between states. In cyberspace, as in all other domains, land, sea, air, and space, these interactions often lead to the emergence of hegemons which are characterised by their predominant influence over global world order and all other states. We examined the emergence and collapse of hegemons in a modelled cyberspace world through the notions of power transition and power diffusion. We used Repast Symphony to construct a simple agent-based model (ABM) of a system of states interacting both competitively and cooperatively in this world. Our simple model parsimoniously captures the character of the real international system of states through simple parameters of *wealth* and *power* determining the outcome of *attack* or *cooperation* amongst pairwise interacting states. We found hegemons of global world order emerged in cyberspace as they do in the other traditional domains from models with these few parameters. And we found that hegemons, contrary to traditional understanding, are not exceptional states but merely occupy the tail of a continuous distribution of power and lifetimes. We also found that hegemony in the system depends on two perhaps unexpected parameters: the difficulty of acquiring power as wealth increases and the amount of cooperation between states. And as a consequence, we argue that cyberspace, as a power-diffuse domain where cooperation is easier than elsewhere, is less suited to the kind of hegemony we see in the traditional domains of state interaction.

1. Introduction

Cyberspace is a power-diffuse domain when compared to the traditional domains of warfare and statecraft, land, sea, air, and space. Scholars have convincingly argued that cyberspace is an offence-dominated domain with low barriers to entry, one that diffuses power away from traditionally powerful states and towards historically marginalised actors [1]. Procuring information communication technology (ICT) systems, acquiring zero-day exploits, and conducting computer network operations (CNOs) are considerably cheaper than large-scale military operations in the domains of land, sea, air, or space. As an example, the Sony Pictures Entertainment hack is a paradigmatic case of a historically marginalised state actor (North Korea) manipulating a third party (Sony) to create preferential outcomes for itself by pursuing aggressive statecraft against the United States, the hegemon of world order [2].

Nye [3] has argued that *power diffusion*, a novel process only now becoming fully appreciated with the emergence

of cyberspace, is one of two critical processes currently influencing world order. The other is *power transition*, a familiar historical process.

The novel process of *power diffusion* can be seen as a weakening of the power differentials between states when we compare the traditional domains of warfare and statecraft, land, sea, air, and space, with the new domain of cyberspace. Cyberpower requires a reevaluation of the power differentials, since the threshold to becoming a relevant actor in cyberspace is lower than in other domains, and the powers of these actors are clustered closer together. As Nye [3] imagines it, if we think that all the states are players on the world stage, then “States will remain the dominant actor on the world stage, but they will find the stage far more crowded and difficult to control” (p. 114).

Analysts of international order must now consider cyberpower in their calculus of factors that affect the rise and fall of states [4], that is, *power transition*. We think that this is particularly so in the important case of global hegemons. Hegemons have disproportionately greater capacity,

and (usually) disproportionately greater success, in creating preferential outcomes for themselves [5]. Every hegemon of global world order is defined by its predominant influence over all other states. Despite their preeminence, hegemonies are not impervious to the anarchical conditions of the international system that permit both cooperation and conflict.

The rise and fall of hegemonies, as for all states, are captured in the process of *power transition*. It is particularly important today. While it is true that the global world order is forever in flux, there is a growing perception amongst scholars that the United States is in relative decline due to the “rise of the rest,” especially since states in the crowded neighbourhood of the Indo-Pacific continue to rise towards great-power status [6].

Because cyberspace is a new and still evolving domain, empirical studies of power diffusion and power transition in cyberspace are not yet available. There have been a number of recent calls from academics, policymakers, and practitioners to better explain, model, and predict the kinds of competitive behaviour we are seeing between states in cyberspace [7]. For example, Kello [8] has called for a synthesis of existing theoretical paradigms with emerging cyber phenomena through the use of computer modelling. We therefore chose to analyse the problem through simulation. We designed a simple agent-based model to capture both the familiar process of transition and the novel process of diffusion in an abstract domain of conflict and cooperation. We simulated the impact of differing degrees of power diffusion on the frequency of power transitions in a domain, and we analysed the resulting emergent hegemonic behaviour. We imagine that the higher levels of power diffusion more closely emulate conditions in cyberspace.

By choosing an empirical simulation approach, we are following in an agent-based modelling tradition grounded in early work by Bremer and Mihalka [9] and later Axelrod [10, 11] and Cederman [12, 13]. These early studies broadly exposed the possibilities of simulation in this field. They were often concerned with the emergence and dissolution of states in quasi-geographic settings (e.g., Cederman and Girardin [14]) and with the balance between cooperation and competition (e.g., Bearce and Fisher [15] and Majeski [16]). The model reported here builds on this work in the particular context of cyberspace, where power diffusion is a key consideration [3] and geography much less so.

As described below, our observations and analysis of the model generated two hypotheses. The first hypothesis is that hegemonies maintain their status more easily in a domain where it is difficult to gain power. Thus, in cyberspace we expect it to be more difficult for any power to maintain hegemony. The fewer resources that need to be invested to achieve a given level of power, that is, the more power diffusion there is in the model world, the more competition a potential hegemon will face. The second hypothesis is that cooperation amongst states leads to more transitions amongst great powers, meaning there is less stability within the international order. When cooperation increases in the system, as it can more easily in cyberspace, we expect to see a greater growth of wealth overall, yet smaller relative power amounts amongst states. Large states end up with more competitors, and *their* positions end up less secure, even though their wealth is higher overall than it otherwise would have been.

2. Description of the Model

We constructed a simple agent-based model of states competing and cooperating with one another using Repast Symphony [17]. We ran our experiments with 20 states because that number is small enough to be manageable and yet large enough to elicit interesting behaviour. (The essential behaviours do not change with 100 states.) Each state has a *power* that represents the state’s probability of winning in a conflict. This is our proxy for the ability of a state to create preferred outcomes, as power is traditionally defined [3]. From a contemporary realist perspective, we assume that states are power maximisers and that power depends on wealth [18]. The *wealth* in our model is intended to be an abstract measure of the total resources at the disposal of a state. The amount of power (the probability of winning a game) purchasable by a given amount of wealth can be seen as a feature of the domain of conflict. In a more power-diffuse domain, it is easier to purchase power.

Each state in the model makes one move per time step. A move for one state consists of a choice of either cooperating with or attacking an interaction partner. The partner is selected at random and also makes the same choice. Thus, with twenty states, twenty bilateral interactions take place per time step. Each state initiates one interaction, and each state is, on average, selected as a partner once per time step.

If both states choose cooperation then a mutually beneficial interaction takes place. The wealth of both states increases marginally. In an environment where the vast majority of states are cooperating, the overall wealth of the system grows. If either state chooses to attack then a conflict occurs. The probability of victory in a conflict is higher for the more powerful state, but there is always a chance the weaker state will prevail. In the model, the victor steals some wealth from the loser, representing the spoils of war. The victor also loses some wealth, which represents the cost of waging war. Thus, lose-lose interactions can take place.

Our model has been constructed with the fewest parameters possible. It is defined by the probability that a state will choose to cooperate and by the power function (the power, or probability of conflict victory, resulting from a given level of wealth). We argue that power diffusion is captured in the power function.

There are, of course, other agents in cyberspace. So-called nonstate actors include large firms, nongovernmental organisations, and criminal syndicates. Their motivations vary from profit and prestige to recreation and revenge. The sophistication of hacking tools and techniques of nonstate actors is highly variable, which affects the amount of power they ultimately wield. We think nonstate actors are a class deserving further investigation, although we confined this first analysis to nation states. Future comparative analyses could explore the relationship between nonstate actors’ motivation, sophistication, and power, which can be contrasted against the same parameters of states, for instance.

3. Tuning the Model for Cyberspace

The model, as described, is both minimalist and generic. It seeks to capture, with very few parameters, the essence of

states interacting through both cooperation and competition, that is, the essence of international affairs. We wish to explore how states interact in cyberspace. We want to see whether and how the forces at play in the traditional domains of statecraft, land, sea, air, and space, might lead, in cyberspace, to the traditional outcomes of statecraft, the traditional international order of one or a few hegemonic or great powers, a handful of middle powers, and a tail of many weaker powers.

To do this, we need to account, in the model, for two canonical properties, geography and power diffusion, that distinguish cyberspace from those traditional domains.

Geography is most apparent in interactions between states in the land domain, and somewhat less so, but still importantly, in the other traditional domains, sea, air, and space. States have always interacted more easily with their geographic neighbours than with states a long way away. The wars throughout history in continental Europe as well as the modern creation of the European Union amongst contiguous neighbours each attest to the power of geography in the land domain. Sea, air, and space power allow states to circumvent geography to an extent, but at a great cost to national wealth. Generally, only great powers, such as the US or Russia, have the wealth and resolve to project power in the sea, air, and space domains that effectively reduces the importance of geography.

But geography has a minimal impact on cyberspace. States may just as easily interact in cyberspace with distant states as with geographic neighbours. Our model in its fundamental form, in contrast with some earlier models of statecraft [9, 14], eschews geography completely, allowing states to interact without regard to any concept of neighbourhood. However, it should be noted that the model framework is flexible enough that explicit geographies could be easily encoded in future work.

Power diffusion is in some ways a more subtle property of statecraft than geography, but its significance, especially in the cyber domain, has been argued powerfully by Nye [3]. We, like Nye, see a weakening of power differentials between states in cyberspace. And as discussed above, we should expect a weakening of the power differentials between states and lower barriers to entry to active participation in the domain.

To incorporate the notion of power diffusion in the model, we need a functional relationship between a state's wealth and its power that reflects these ideas. As discussed above, we see wealth as an abstract measure of the resources available to a state, a quantity that can grow or shrink as a result of interactions with other states. And we see power as a quantity that measures the resources a state can apply to war or competition. If states go to war, the probability that a state will win depends on the powers of the two states, with the more powerful state having a higher probability of winning (see Appendix A).

We argue, with Schmidt [18], that a state's power depends on its wealth, but that the relationship is nonlinear. It needs to account for not only differences in the degree of power diffusion across the different domains of statecraft, but also differences between states. We know, for example, that there exist states like Japan which, though wealthy, only weakly convert that wealth to power, and states like North Korea

which, though poor, strongly convert what wealth they have to power.

We describe below (Section 4) some general classes of power functions that could capture these relationships, and then (Section 5) we describe our strategy for exploring their effects on statecraft in the cyber domain.

4. Power Functions: Creating State Power from Wealth

All the domains where states interact vary on several parameters, amongst which are the degree of connectivity between states and the degree of power diffusion in the domain. This variation accounts for differences in the behaviour of states amongst domains. In our model, we focus on power diffusion and assume, for simplicity, that all states can interact with all other states, mimicking the high connectivity of cyberspace. We identify the degree of power diffusion with the difficulty of gaining power and model this parameter with a power function, which provides the conversion rate from wealth to power. The shape of the power function determines the outcomes of conflict or cooperation amongst states and thus the overall behaviour of the system. Figure 1 shows some general classes of power functions.

Figure 1(a) is a function that curves rapidly upward. Thus, small increases in wealth produce large increases in power, which makes it easier for a large state to dominate smaller states and get rapidly more powerful. In a domain with this extreme power function, there is no limit to the power of a state. We include this class for completeness only. We imagine that such a function would only exist in highly asymmetric cases, perhaps advanced colonial powers against primitive indigenous peoples.

Figure 1(b) is a function that curves downward, representing diminishing returns on power. This creates a natural soft limit on how powerful a state can realistically get. We think asymptotic classes like this one reflect the common reality of the conversion of wealth to power. This class would cover many of the power relations in historic domains such as land and sea.

In Figures 1(c) and 1(d), the functions are sigmoidal. They represent the more general case where an initial investment produces an initial burst of power, followed by gradually diminishing returns. This class of functions is characterised by the width on the x -axis of the upward-curving region. States in this region are much weaker and easily dominated by the stronger states outside of it. If the region is wide, like in Figure 1(c), it takes a lot of wealth before a state reaches the soft power limit. If it is narrow, like in Figure 1(d), then it is easy for states to reach a high level of power. Thus, in Figure 1(d) the power function represents a power-diffuse domain like cyberspace, while in Figure 1(c) the function could stand for expensive technology-rich domains, like air and space, where power is harder to come by.

Note that in Figure 1(a) the function could be seen as having an infinitely wide *weak region*, so that *all* states are easily dominated by those larger than them, while in Figure 1(b) the function does not have such a region, and so increasing wealth never allows a state to be utterly dominant.

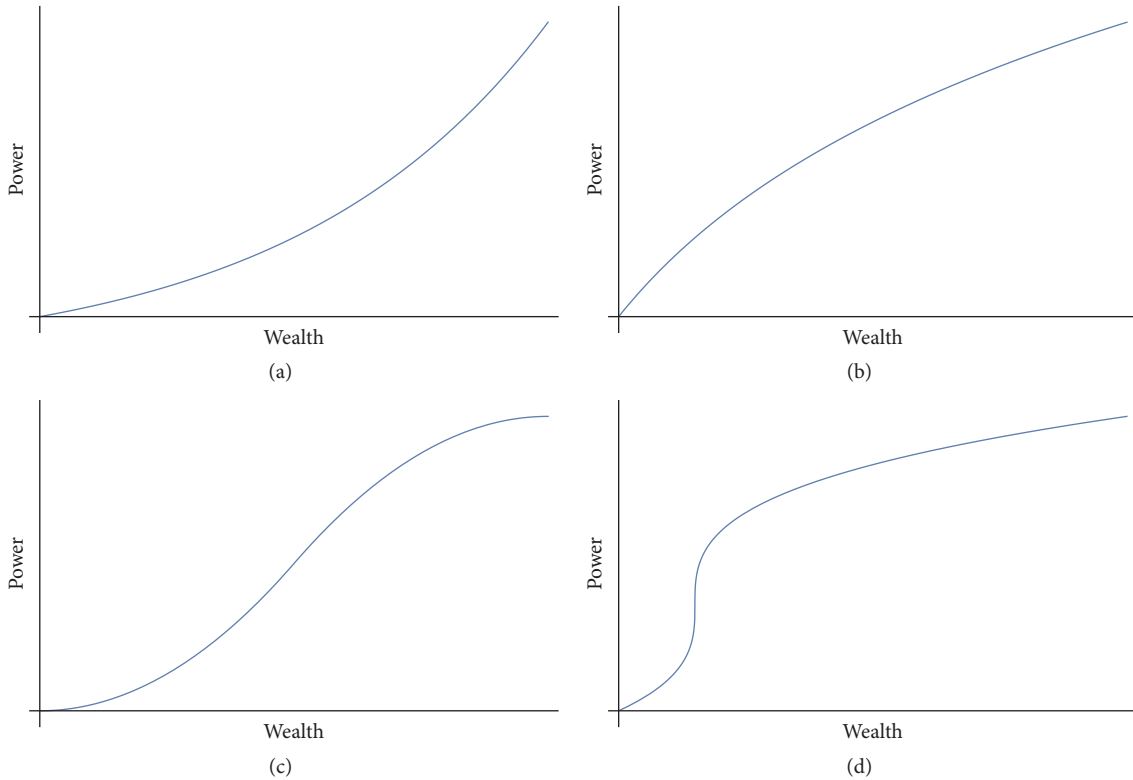


FIGURE 1: Four possible power functions, representing (a) limitless power, (b) a soft limit to power, perhaps as in the historic domains of land and sea, (c) perhaps an expensive technology-rich domain, like air and space, and (d) a power-diffuse domain like cyberspace.

5. Exploration Strategy

We made two sets of experiments to explore the behaviour of our model. The first explored the extent to which the model could capture the traditional dynamics of the international order in traditional domains (Section 5.1). The second built on the understanding gained in the first experiments to explore the dynamics in the cyber domain (Section 5.2).

We experimented with two parameters, the probability of cooperation between states and the degree of power diffusion. The former was implemented directly as a factor in the behaviour of the agents, while the latter was varied indirectly by varying the power function.

For the first set of experiments, we used, for simplicity, a simple *linear* power function (which does not have a power-diffusing weak region as in Figure 1). And then, for the second set of experiments, we used a *logarithmic* power function (which allowed the degree of power diffusion to be varied across the range implicit in Figures 1(c) and 1(d)). In each set, we studied various levels of cooperation.

5.1. Behaviour of States in Traditional Domains. Figure 2 illustrates the simple linear power function we used in these experiments, and Figure 3 shows the typical behaviour of the model as states' probability of cooperating is varied from low to high. We show the evolution from tick (time step) 1000 to tick 2000 because it takes about 1000 ticks for the model to settle down into a stable state, after which the behaviour does




FIGURE 2: The linear power function starts out flat until it reaches what we call the minimum wealth, which is set to the value of 1.0. States below this wealth are unable to fight wars with larger ones, so this power function restricts the “major players” in the game to a subset of all the states.

not change significantly, even beyond 100,000 ticks. Note that the vertical scale varies between the plots.

Figure 3 shows that, with a low probability of cooperation of 0.3, most states remain in the region below the minimum wealth, where they cannot have conflicts. Whenever they rise above, they are quickly put down by the power of the large states. Furthermore, the correlation between the red and the yellow state reveals that the most meaningful interactions

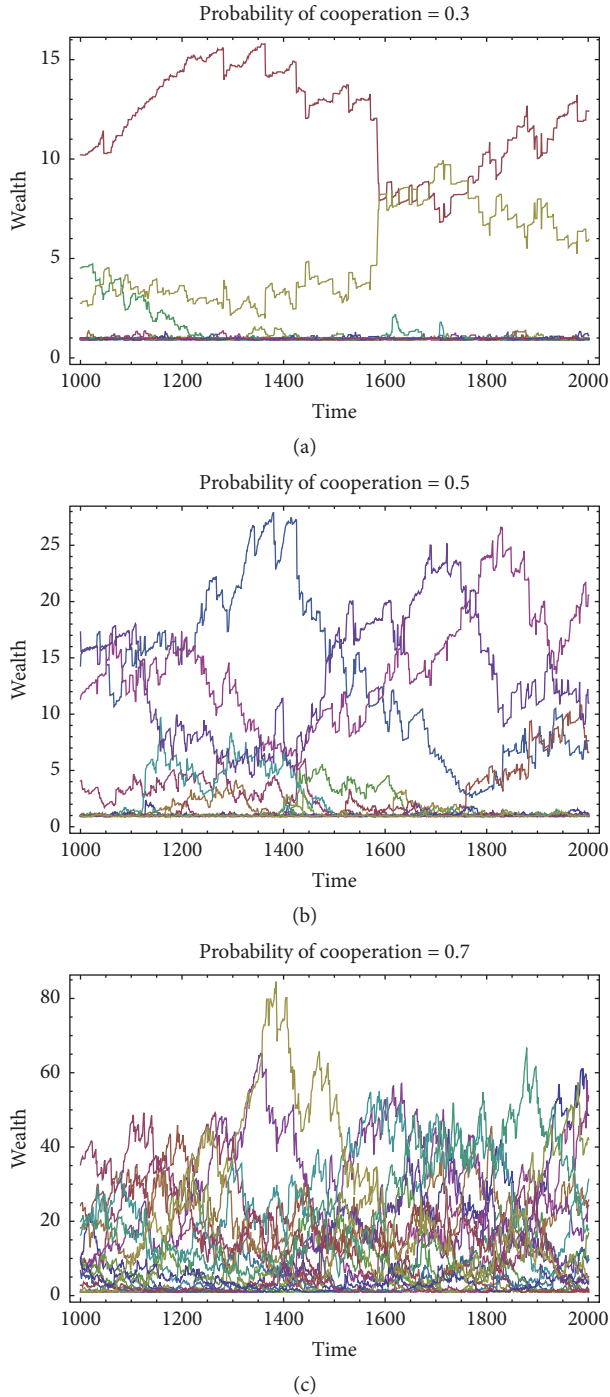


FIGURE 3: Typical behaviour with the linear power function, across several values of the *probability of cooperation* parameter. Each colour represents one of the 20 states.

take place between two large states, not between large and small.

At a cooperation probability of 0.5, small states are able to rise above the minimum wealth by cooperating with larger ones. This cooperation allows the upper layer (above minimum wealth) to hold more states, but also makes it more dynamic: states enter and exit the layer more frequently.

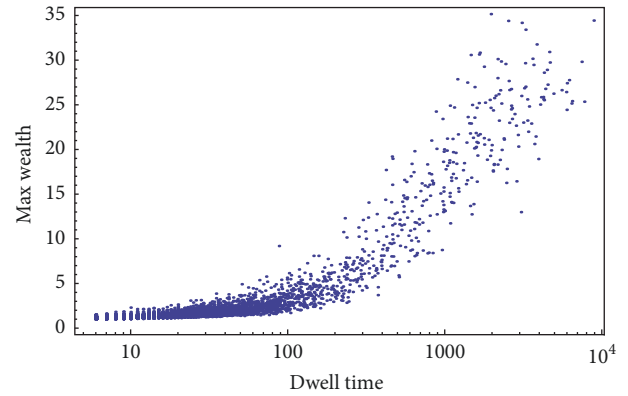


FIGURE 4: For a probability of cooperation of 0.5, we plot a point each time a state rises above the minimum wealth. The *dwell time* is the time spent above the minimum wealth, and the *maximum wealth* is the maximum wealth achieved during that dwell time. A logarithmic scale is used on the horizontal axis for clarity. The smoothness of the distribution indicates that hegemonic states (in the top right of the plot) are not special.

With a high probability of cooperation of 0.7, the dynamism increases further. This happens because as the number of large states increases, so does the probability of a large defeat for any particular state, so none remain dominant for long.

In summary, more cooperation means fewer wars, so the equilibrium capacity of the upper layer increases. Increased dynamism results from each state having more opportunities for wars with other large states, which are the only kind of war that can significantly affect a large state's wealth with this power function. To demonstrate this more rigorously, we have Figures 4 and 5.

Figure 4 shows the distribution of state wealths and lifetimes over the course of a single model run. (We chose the awkward construction “wealths” because a nation state can achieve any number of states of wealth, and we need to distinguish these individual values and treat them as a set. But we remain unhappy with this corruption of the English language.) It shows that hegemons (such as the dominant red state in Figure 3(a)) are not produced by a special process, but are simply normal states at the extreme end of a range of different levels of power.

Figure 5 combines data from model runs at several different probabilities of cooperation. It demonstrates that increased cooperation allows the upper layer (above minimum wealth) to hold more states, but the states that are there do not last quite as long.

We can also see the effect of cooperation by measuring how much time states spend possessing different amounts of wealth. We ran the model 100 times with different probabilities of cooperation. A few pilot runs established the typical range of wealths that states achieve. We split this range into buckets and calculated the number of states whose wealth was within each bucket each tick. Summing these values across all the ticks for a single model run produces one vertical slice of Figure 6. The horizontal axis allows the results for different probabilities of cooperation to be compared.

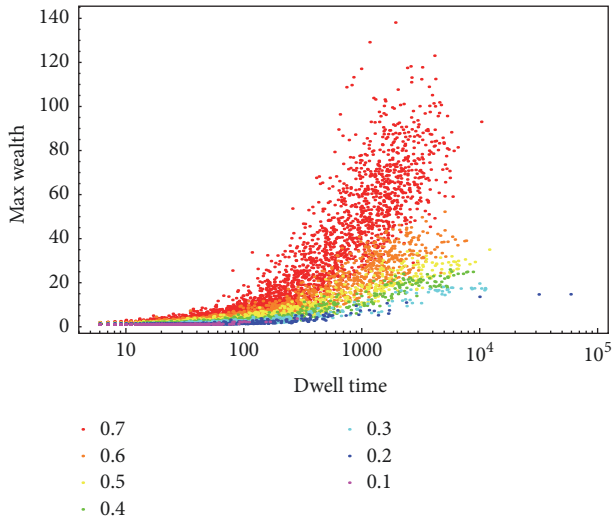


FIGURE 5: As for Figure 4, but with several data sets of different probabilities of cooperation now combined. The colours (see the keys of the figure) range from red for high probability of cooperation to blue and purple for low. At a higher level of cooperation, the maximum wealth achieved by states is much larger on average, but they do not last quite as long as the longest-lived hegemon in the bottom-right when cooperation is low.

Figure 6 is coloured to show the frequency of different levels of wealth achieved for different probabilities of cooperation. Thus warm colours indicate that states occupied that level of wealth at that probability of cooperation more often, while cool colours, less often.

Consider, for example, the vertical slice in Figure 6 at a probability of cooperation of 0.4. Moving upwards, we see first a thin green layer, states just above minimum wealth, then a thick blue band, and another green band. This suggests that, in that model run, states were often either just above minimum wealth or far above (in the second green band), but typically skipped right past the space in between. The point at which the colour turns white indicates the highest wealth achieved by any state in that model run.

Figure 6 shows the maximum wealth increasing with higher probability of cooperation, although it is not achieved very often. It also shows that states tend to concentrate in a band somewhat above the minimum wealth level, but that this band spreads out and encompasses most of the states at higher levels of cooperation. The bright spot on the left corresponds to the one or two hegemon that remain in control for the entire model run at that low level of cooperation. The red region on the right indicates that most states are able to achieve a higher level of wealth when there is more cooperation. This is again the same observation that more cooperation leads to increased capacity and dynamism.

When we ran the model with more states, the results were effectively the same, although plots like those in Figure 3 would appear compressed, as if the model were running faster. This makes sense, since with more states there are more interactions happening per time step, and hence each state engages in more cooperation and more conflicts.




FIGURE 6: The distribution of state wealth for different probabilities of cooperation. Each vertical line represents one model run at a particular probability of cooperation. Warmer colours indicate that states occupied that level of wealth more often. Wealths below the minimum, indicating inactive states, have been truncated.

The results above show that, with the linear power function, single hegemon invariably emerge and that increasing cooperation leads to higher wealth overall, but reduces the dwell time of such hegemon. These are satisfying results because they closely emulate the observed international order seen throughout history in the traditional domains of land and sea. The emergence of relatively long-lived hegemon when there are low levels of cooperation (Figure 3(a)) looks very similar to the European powers at the height of their mercantilist strategies in the mid-19th century. And the world with higher levels of cooperation and wealth (Figure 3(c)) looks not unlike the postwar Bretton Woods world, with rising wealth in middle powers.

The results also show a clear power-law relationship between maximum wealth and dwell time of states (Figures 4 and 5). While we know of no empirical studies of power transition in traditional domains, it is nonetheless satisfying to see that a simple model driven by competition and cooperation leads to the emergence of power-law relationships like those seen empirically in other complex systems with simple opposing internal forces. See, for example, Bettencourt [19] on city size, Cederman [20] on war, and Hatton et al. [21] on predator-prey relationships.

We conclude that the model, using a simple linear power function, captures the essential behaviours of power transitions in the traditional domains. This allows us to move with some confidence to explore the behaviour of states in cyberspace.

5.2. Behaviour of States in Cyberspace. Figure 7 shows the logarithmic power function that we used in these experiments.




FIGURE 7: The logarithmic power function $p(w) = \ln^k(1+w)$ for two values of the exponent k . It has two regions, one that curves up and one that curves down. States in the second region have an advantage over those in the first region. In other words, once a state has enough wealth to reach the steep increase in power, it joins an elite club of powerful states. The less wealth required to reach that region, the more power-diffuse a domain is. Reducing the exponent shrinks that region and so models increasingly power-diffuse domains.

This function is useful because it has a parameter, the *power exponent*, that can be varied to alter how easy it is to acquire power and hence the degree of power diffusion. A higher exponent corresponds to less power diffusion. Thus we can emulate the qualitative power functions of Figure 1 by varying the exponent to create a broad class of function shapes that encompass both shapes of Figures 1(c) and 1(d) in their spectrum. This is described in more detail in Appendix C.

Figure 8 shows the behaviour of the model as cooperation increases. The power exponent is set at 1.75 for these experiments. This creates a more or less linear power function, and the results are broadly similar to Figure 3, indicating that our analysis of the effects of cooperation applies in this case as well. It also shows that the notion of power diffusion can be effectively added to the model to provide a smooth path from the dynamics in traditional domains to the new domain of cyberspace.

Figure 9 shows the changes as the power exponent is varied. As the power exponent drops (and hence power diffusion increases), the world moves from a single hegemon, to a cluster of great powers, to a free-for-all.

Figure 10 is comparable to Figure 4 and demonstrates that the overall behaviour of the model with a logarithmic power function is similar to its behaviour with the linear power function.

Figure 11 allows us to analyse the effect of changing the power exponent (and hence the degree of power diffusion). We see that less power diffusion produces more and longer lived hegemonic states, but at the expense of the rest. The green dots that are not hegemons (the hegemons are in the top right) are concentrated in the bottom left, beneath the bulk of the red dots from the more power-diffuse domain. Less power diffusion increases the difficulty of acquiring a significant amount of power, reducing the total number of large

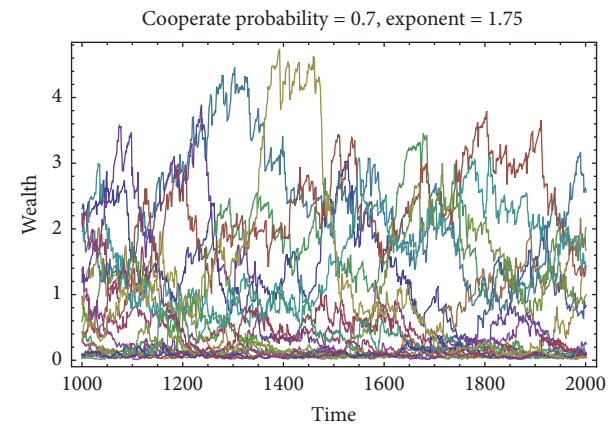

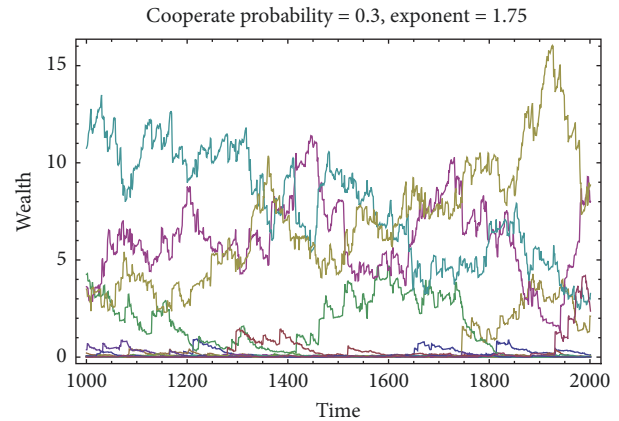


FIGURE 8: The typical behaviour of the model with the logarithmic power function, with an average level of power diffusion, as probability of cooperation is varied. Each colour represents one of the 20 states in the model.

states and the competition for the dominant spot, allowing hegemons to last for longer.

Finally, Figure 12, read together with Figure 11, clearly shows a transition in the dynamics of the whole ensemble of states as power diffusion increases (and thus the exponent decreases). Figure 12 shows the results of several hundred runs of the model with different exponents. It is coloured like Figure 6 to show the frequency of different levels of wealth achieved but, this time, for different levels of power diffusion

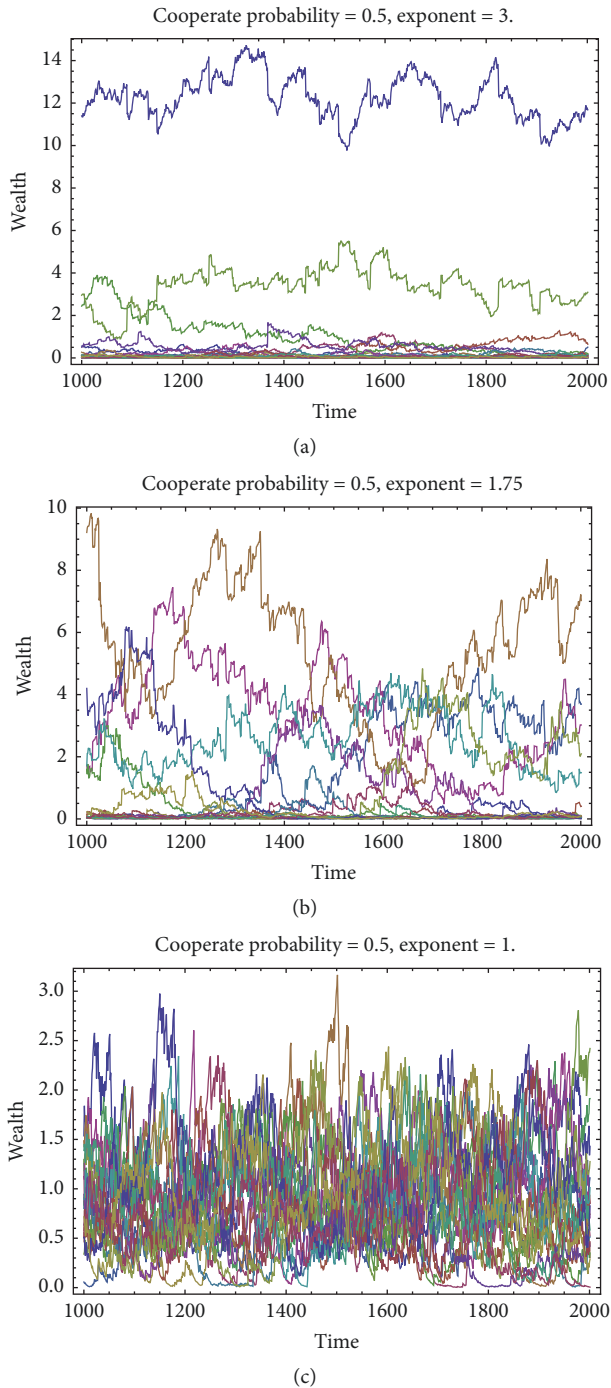


FIGURE 9: The typical behaviour of the model as the power exponent is varied. (a) represents the least power-diffuse domain.

(as measured by the power exponent). Thus, warm colours indicate that many states frequently occupied a level of wealth at that level of power diffusion more often, while cool colours represent the inverse.

At the left of Figure 12, power diffusion is low and the exponent high. A vertical slice in this region shows, moving downwards in wealth, a few (at the limit, one) wealthy, powerful, long-lived (see Figure 11) states, hegemons, that suppress




FIGURE 10: The shape of the distribution of maximum wealth and lifetimes using a logarithmic power function with an exponent of 1.75 and a cooperation probability of 0.5 is similar to that seen in Figure 3.




FIGURE 11: Data from nine runs with different power exponents and a fixed cooperation of 0.5 is plotted here. Each dot represents one growth and decline period for one country in a particular run and plots the maximum wealth achieved during that time. Redder dots correspond to a more power-diffuse domain than the greener dots.

the wealth and power of the states below them. Thus immediately below the hegemons is a relatively sparse region of a few persistent middle powers (in green and yellow). Most states (thus, in warmer colours) are below them at low levels of wealth. Below these poorer states are a few states (thus, in cooler colours) at the bottom that are very poor.

As power diffusion increases towards the right of Figure 12, the ensemble of states is becoming progressively wealthier as the floor of minimum wealth grows (coloured white), but the wealth of the wealthiest states is declining. Taking a vertical slice in the middle of the figure and reading downwards, we see a transition beginning as a widening band of abundant reasonably wealthy states (thus, a widening band of yellow progressing to red) emerges immediately below the




FIGURE 12: The distribution of state wealth for different power exponents (a lower exponent means more power diffusion). White indicates no states, and warmer colours indicate that states occupied that level of wealth more often. The horizontal axis has been reversed to emphasise the increase in power diffusion from left to right.

wealthiest states. The wealthiest states are now cheek by jowl with a growing number of states that are nearly as wealthy. That is, the wealthiest states are no longer suppressing those immediately below them. The uppermost layers of wealth now progressively hold more states. Below that emerging band, we see the band of states with low levels of wealth gradually dissipating as wealth is spread more widely.

Compared to the clearly hegemonic states at the top left of Figure 12 (and top right of Figure 11), the wealthiest states in the middle of Figure 12 are barely hegemonic. They are relatively poorer and shorter-lived states (see Figure 11). They are unable to suppress the wealth and power of most of the other states which now cluster just below them in wealth and longevity in an ensemble with far more flux (Figure 9).

At the extreme right of Figure 12, most states now have more wealth (the warm band towards the top right is now fully red), but the wealthiest states are now not so wealthy. And we know, from Figure 11, that these wealthiest states do not last as long as the wealthiest states in the less power-diffuse regions further to the left. We no longer observe hegemons, but rather a set of states whose wealth and power are constantly churning (Figure 9).

These results show that single hegemons once again invariably emerge and that they are more stable when there is less power diffusion, that is, when the power function was more like Figure 1(c) rather than Figure 1(d). They also show, as did the results in Section 5.1, a clear power-law relationship between maximum wealth and dwell time of states.

6. Discussion

We have produced a simple model of state interactions. Clearly it lacks any of the details of actual interactions, but the simplicity of the model means that we make very few

assumptions, and thus the broad conclusions are widely applicable [22].

We defined wealth as an abstract measurable value of the total resources at the disposal of a state. This concept contains metrics of military, economic, geographical, and scientific indicators that constitute national wealth. In the real world, qualifying what resources or attributes constitute national wealth is fraught with inconsistencies and it has confounded analysts in the past. However, resolving these questions is not necessary to experiment with a simple model like our own and glean useful insights from it.

In cyberspace, wealth can buy zero-day exploits. Euphemistically, these are “guns with only one bullet” [23]. Wealth can procure computer systems. Wealth can employ ICT security architects to design safe systems. And wealth can “grease the palms” (or fingertips) of black-hat hackers [24]. Wealth, as we conceptualise it, is an input that states can utilise to purchase the bits and bytes necessary to defend against or carry out statecraft in cyberspace.

In contrast, power is the ability to perform statecraft within the environment of cyberspace to create preferential outcomes. Power is the timely deployment of purchased exploits towards strategic ends [25]. Power can create preferred outcomes, shape agendas, and create structural frameworks.

For the purposes of our model, we defined power as a quantity determined by the wealth of a state, with the exact relationship depending on the domain. A more power-diffuse domain, like cyberspace, allows a state to gain more power more easily, with a smaller investment of wealth. The conversion from a given quantity of wealth to a particular amount of power can be varied. This power function represents the real-world capacity of states to turn resources into outcomes. We argued above that different power functions allow us to model the amount of power diffusion.

We have identified two factors that affect the distribution of power under these simple dynamics. One is the rate of cooperation. Like the conflicts, the cooperation in our model can stand for many real-world interactions, such as trade, technological diffusion, and intelligence sharing. We have demonstrated that, with more cooperation, all or most of the actors are better off than they otherwise would have been, but the more hegemonic ones find it more difficult to hold onto power for as long. Historically, this suggests that hegemons who share power and technology with allies are likely to be better off, but also less clearly dominant, than those who prevent any competition by force. A different example of a power structure that could also be modelled this way is the ruling party or family in an authoritarian state. By ensuring complete dominance, the rulers secure their position, at the cost of the prosperity of their nation [26]. With more liberal policies, everyone is often better off, and the wealthiest people may be wealthier, but this would not be in the interests of the existing elite, since the ability for a specific single actor to maintain power in a cooperative environment is compromised.

The other significant factor was the shape of the power function, which is how different domains of state interaction can be abstractly represented in our model. A function that

starts off flat, rises sharply, and then flattens again represents a domain where the small players cannot significantly affect the bigger players, but the biggest players are comparable to each other. This, we claim, is the shape of the power distribution in the traditional domains, and our model shows that it produces hegemonic behaviour with only occasional upsets. Increasing the size of the early flatter region corresponds to a greater barrier to entry, and our model correctly predicts long-term stable hegemony in this case. Finally, reducing the size of the region (a smaller exponent) means lowering the barriers to entry, for example, in cyberspace, where although the power of a nation state is not easy to come by, small actors can nonetheless have a large impact [27], and our model predicts much greater instability in this power-diffuse domain, in line with the expectations of some scholars [28].

7. Conclusion

We have presented a simple model of competition and cooperation between states. Two key factors that affect the emergence of hegemony are the amount of cooperation and the ease of acquiring power. Less cooperation and greater barriers to power favour the establishment and maintenance of hegemonic actors. We argue that cyberspace as a domain of conflict between states is less suited to hegemony on both of these metrics than traditional domains like land, sea, and air. We therefore agree with Nye [3] that cyberspace, as a more power-diffuse domain, is unlikely to support the kind of hegemonic behaviour we have seen states engage in in the other domains.

8. Future Work

We now intend to evolve a more sophisticated model from this fundamental framework of abstract states in abstract domains. We plan to carefully parameterise the states, in a stepwise manner, to give them some of the canonical characters and behaviours of the set of real-world states. And we will likewise parameterise the domains to create more realistic topologies.

Appendix

A. Model Details

An agent-based model was constructed using Repast Symphony. The model's source code, including instructions for how to get it running, is available at <https://www.openabm.org/model/4600/>.

A.1. Parameters. The complete list of parameters (and default settings) is as follows:

- (i) Number of states, $N = 20$
- (ii) Linear or logarithmic power = linear
- (iii) Minimum wealth (linear only), $\text{MIN} = 1.0$
- (iv) Power exponent (logarithmic only), $k = 1.75$

- (v) Probability of cooperation, $c = 0.5$
- (vi) Wealth modification factor, $f = 0.01$
- (vii) War steal factor, $s = 0.1$
- (viii) Cooperate penalty, $\gamma = 0.5$
- (ix) End at tick = 10000.

Throughout the text, the symbols above will be used to refer to these parameters, and wherever the text does not explicitly mention them, they were at their default values.

All states begin at MIN wealth (for linear power) or 1.0 wealth (for logarithmic power).

A.2. Model Dynamics

- (1) At every time step (tick), every state gets a turn to select another state and interact, in a random order. N interactions happen per tick, and each state interacts twice on average.
- (2) Each state in an interaction chooses a move, to either attack (ATT) or cooperate (COOP). They make this choice randomly according to the probability of cooperation c . If a war occurs, the probability of a state being victorious depends on the powers of the two states. Their wealth is altered depending on the outcome. See below for details.
- (3) For logarithmic power functions only, all states' wealths are altered in between ticks according to $w_{t+1} = w_t / (1 - f + 2fp^2)^2$. We refer to this as *growth correction*, and it exists to counteract the otherwise inevitable decline produced by that function.

A.3. Power Functions. We shall refer to the wealth of arbitrary states as w , w_a , w_b and their power as p , p_a , p_b . Power is a function of wealth, and, based on some preliminary exploration not recorded here, we selected two power functions of interest:

- (i) Linear power: $p(w) = \max(w - \text{MIN}, 0)$
- (ii) Logarithmic power: $p(w) = \ln^k(1 + w)$.

A power of zero (as happens below MIN for the linear function) has the special effect of causing a state to ignore wars. Furthermore, as analysed further below, the concave up region in the logarithmic power for $k > 1$ is a region where states are weaker and tend to lose wealth to those above. Thus, both power functions contain a region of low wealth where states present less of a threat to wealthier ones. As we observe later, this correspondence is responsible for similar behaviour under both functions.

We chose these two functions for simplicity and utility: they allow us to explore the qualitative dynamics of power as a function of wealth in our unparameterised model.

Our linear power function allows for a power of zero at some presumably low wealth (to allow for the case of a state ignoring wars) and then shows an unbounded linear relationship between power and wealth. It is essentially a base case to compare with richer relationships.

Our logarithmic power function addresses the issue that many “growth” functions, especially those involving living things, have a sigmoid shape, starting slowly, increasing, and then slowing down. In ecology, these often explicitly asymptote to a “carrying capacity.” Our logarithmic function has a similar sigmoid shape, but no upper bound, since we do not wish to constrain the growth of power in this way. The growth just slows exponentially over time. Thus, using this function allows us to explore the qualitative dynamics of a range of different power-wealth relationships simply by varying the exponent (see Appendix C).

A.4. Interactions. If states a and b both cooperate, they each gain $f(w_a + w_b)/2$ wealth.

Otherwise, a war occurs, and state a wins with probability $p_a/(p_a + p_b)$. If one state cooperated, their power is multiplied by γ for this calculation. If b loses, w_b 's wealth is transferred to a . Finally, the winner of the conflict has to pay a war cost of $f(w_a + w_b)$, which may result in overall losses instead of gains for the victor. If one state has a power of zero in this interaction, then no war occurs and all wealths remain the same as before. Effectively, states with wealth below MIN are insignificant compared to larger states and do not engage in wars. They cooperate as normal.

A.5. Growth Correction. For the logarithmic function (or indeed any power function that is always positive), the overall growth and decline do not depend on the shape of the function. Consider the average gain or loss for a single state in a single interaction. If we assume (to simplify calculations) that all states have equal wealth, the state has c^2 probability of gaining fw wealth and $1 - c^2$ probability of being in a war. When talking about overall growth, we do not care about the transfer of wealth between states, only about wealth that is removed from the system, war costs. Since only one out of the two states will have to pay war costs, the average loss is $2fw$ with probability $(1 - c^2)/2$.

Thus, on average, a state with wealth w will have wealth $c^2w(1 + f) - (1 - c^2)w(1 - f)$ after an interaction, which corresponds to multiplying w by $1 - f + 2fc^2$. Since each state interacts twice on average, a correction multiplier of $1/(1 - f + 2fc^2)^2$ has to be applied every tick to avoid uncontrolled (and meaningless) growth or decline. This multiplier may be greater or less than one.

For the linear power function, the more states have a wealth below MIN, the fewer wars there are, while cooperations are unaffected. Thus, the more states drop below that line, the more growth everyone experiences. If the parameters are such that the states experience overall decline, they will drop until enough states have crossed that line, and the decline ceases. The linear power function is therefore self-stabilising for $c < 1/\sqrt{2} \approx 0.71$. Above that value, corresponding to the point at which the growth correction would have to be smaller than one, states tend to gain from each interaction on average, even when everyone is above the line. Thus, the overall wealth grows uncontrollably. For this reason, Figure 6 is cut off at $c = 0.7$.

B. Assumptions

The model as described above inevitably includes some assumptions that may not perfectly reflect the behaviour of states in the real world. Our goal was to minimise the number of assumptions (to minimise the ways in which the model can be wrong) and to make justifiable choices for the remainder.

Many features of real-world states are not reflected in this model; for example, the states cannot communicate and form alliances, they do not have any internal factions, and they do not have any sort of memory. These are all possible ways to expand the model in the future, but adding any one such feature also comes at a cost; it introduces many more assumptions, any of which can fail.

We explained our choice of power functions above: a simple linear function and an unbounded S-shape. In the real world, two states with equal wealth could have different military strength, but we have abstracted away this detail. As an example of our thought process, we do not believe that adding an extra *proportion of wealth used for military* parameter would add any new insights to the model. Instead, we imagine that two such states would be represented in our model as two states with different amounts of wealth.

The possible interactions were intended to fit into a 2×2 payoff matrix (the simplest possible when two states both make a choice), and reflect the basic idea that *cooperation is profitable and war is costly (but the victor might recoup the cost)*. The remaining parameters were either varied experimentally as we describe in the paper, or chosen arbitrarily to be within a range where the behaviour is *interesting* (i.e., where the states do something other than grow/shrink without bound, or have one state dominate everyone else forever). With these constraints, most of the parameters are either arbitrary (e.g., f , which sets the typical scale of the interactions but does not affect their dynamics) or forced (e.g., the ratio between cooperation gains and war costs, which must be 1 on average to prevent unbounded growth).

C. Power Function Dynamics

Given a war between states a and b , with $w_a > w_b$ and ignoring war costs, the expected value for state a is

$$\begin{aligned} E(a) &= -E(b) = p_a w_b s - p_b w_a s \\ &= \left(\frac{s}{(p_a + p_b)} \right) \times (p_a w_b - p_b w_a). \end{aligned} \quad (\text{C.1})$$

The fraction is just a constant factor, but the $p_a w_b - p_b w_a$ term is a very important driver of the dynamics of the model. If this term is positive, the wealthier state stands to gain from the war on average, while if the term is negative, it stands to lose. For a particular pair $a > b$, a positive $E(a)$ tends to increase the wealth disparity, while a negative term tends to decrease it.

For a linear power function $p(a) = ra$, this term is always zero. Hence, with linear wealth (in the absence of other important effects), being wealthier gives no advantage or disadvantage overall relative to poorer states. Other power




FIGURE 13: $k = 1$.

FIGURE 15: Linear power.

FIGURE 14: $k = 25$.

function shapes alter this behaviour. Specifically, the term is positive if a line on a power-wealth graph drawn from the origin to the point $(a, p(a))$ passes above the point $(b, p(b))$ and negative if it passes below. To illustrate this, consider, in Figures 13 and 14, the logarithmic power function for two different values of k .

Adding an exponent greater than 1 has the largest effect near the origin, where it produces a concave up region.

Now suppose a has wealth around 5.0. We plot a straight line through that point and the origin. If b has a wealth in the red region below, $E(a)$ would be negative. In the blue region, it would be positive. At the crossover point, $E(a) = 0$. Note that if the power function was linear, it would lie entirely on that line, so $E(a) = 0$ everywhere.

Effectively, concave up regions of the power function produce greater and greater advantages for increasing wealth, while concave down regions penalize wealth. Note that a disadvantage for the wealthier state is equivalent to an advantage for the poorer state. Also note that the crossover point will shift as a 's wealth shifts, so the concavity of the function is only a guide.

When the exponent is raised, the initial part of the graph becomes concave up; as a result, for small wealth, growth becomes an advantage and we see more hegemonies forming. At higher wealth, however, the advantage disappears, preventing extreme breakaways.

Because of the flat region in the linear power function (see Figure 15), larger states always have an advantage over smaller ones, though the advantage diminishes with increasing wealth.

D. Measurement Technique

To produce the diagrams in this paper, we collected data within Repast Symphony, using the *File Sink* mechanism to write it to files. The resulting files were in an approximately CSV (comma-separated values) format. We used a text editor's search and replace function to clean up some artifacts of the Repast output and produce a valid CSV file; then Microsoft Excel was used to sort the data by the parameter of interest. We have also provided a python script that replicates our text editor and Excel workflow. A batch run in Repast produces both an output file (where each row has a *run number* and corresponds to a single run) and a parameter map, which associates run numbers with parameter settings. We used Excel to combine the two, which permitted sorting. Once a sorted CSV file of the results was produced, we imported it into Mathematica and used Mathematica functions to process the data (including combining data from multiple runs with the same parameters) and to produce the diagrams. The Mathematica 9 notebooks we used are available at <https://www.openabm.org/model/4600/> along with the model code.

To produce the scatter diagrams of maximum wealth against dwell time, we set up each state to start a timer whenever it went above MIN wealth. For the logarithmic power function, we instead used the value $10^{-0.5}$, which was chosen because it seemed to be a natural dividing line on the wealth density plot. The resulting output file was a long list of time/wealth pairs. We also made sure to forcibly stop the timers and collect data from states that were above the threshold when the model run ends. Otherwise, a state that stayed above the threshold the entire time (a very significant event) would be missed. The data collected this way was output separately, so that we could judge whether it distorted the rest or not.

We measured wealth density by setting up a predefined collection of buckets at the start of each run. The number and

size of the buckets are determined within the code itself. Each “bucket” is a single integer, storing how many times a state has had a wealth within a certain range. During the run, each time step, for each state, the bucket corresponding to that state’s current wealth is incremented by one. At the end of a run, the list of buckets is output. A batch run thus produces a list of lists, which Mathematica can render as a *ListDensityPlot*. The data from multiple runs with the same parameters (but different random seeds) can be easily combined by averaging as long as they used the same configuration of buckets.

We used buckets with a logarithmic size (wider at higher wealth levels) because all the relevant actions in the model are multiplicative. Thus, besides the effects of the power function, the behaviour of the model at higher wealth is a scaled-up version of its behaviour at a lower wealth. Using a bucket size that increases in proportion thus more accurately captures the model’s dynamics.

There cannot be an infinite number of buckets, and the wealth of states can grow without bound under some choices of parameter settings. To address this issue, we ran several test runs at the extremes of the parameters we wanted to investigate, to find out the highest and lowest wealth typically reached by the states. The maximum and minimum buckets were then set accordingly.

Conflicts of Interest

The authors declare that there are no conflicts of interest regarding the publication of this article.

References

- [1] P. Cornish, D. Livingstone, D. Clemente, and C. Yorke, *On cyber warfare*, in *Chatham House Report*, London, Chatham House, 2010.
- [2] N. Inkster, “Cyber Attacks in La-La Land,” *Survival*, vol. 57, no. 1, pp. 105–116, 2015.
- [3] J. S. Nye, *The future of power*, Public Affairs, New York, 2011.
- [4] “US-China relations in strategic domains,” in *The National Bureau of Asian Research*, Seattle, Wash, USA, 2016.
- [5] S. G. Brooks, “Can we identify a benevolent hegemon?” *Cambridge Review of International Affairs*, vol. 25, no. 1, pp. 27–38, 2012.
- [6] V. Jackson, “The rise and persistence of strategic hedging across Asia: A system-level analysis,” in *US Alliances and Partnerships at the Center of Global Power*, A. J. Tellis, A. M. Denmark and, and G. Chaffin, Eds., pp. 317–342, National Bureau of Asian Research, Washington, CD, USA, 2014.
- [7] C. Vaishnav, N. Choucri, and D. Clark, “Cyber international relations as an integrated system,” *Environment Systems and Decisions*, vol. 33, no. 4, pp. 561–576, 2013.
- [8] L. Kello, “The meaning of the cyber revolution: Perils to theory and statecraft,” *International Security*, vol. 38, no. 2, pp. 7–40, 2013.
- [9] S. A. Bremer and M. Mihalka, “Machiavelli in machina: Or politics among hexagons,” in *Problems of world modeling: Political and social implications*, K. W. Deutsch, B. Fritsch, H. Jaguaribe, and A. S. Markovits, Eds., pp. 303–337, Ballinger, Cambridge, MA, USA, 1977.
- [10] R. Axelrod, *The evolution of cooperation*, Basic Books, New York, NY, USA, 1984.
- [11] R. Axelrod, “A model for the emergence of new political actors,” in *Artificial societies: The computer simulation of social life*, N. Gilbert and R. Conte, Eds., pp. 15–32, University College Press, London, UK, 1995.
- [12] L.-E. Cederman, *Emergent actors in world politics: How states and nations develop and dissolve*, Princeton University Press, Princeton, NJ, USA, 1997.
- [13] L.-E. Cederman, “Computational models of social forms: Advancing generative process theory,” *American Journal of Sociology*, vol. 110, no. 4, pp. 864–893, 2005.
- [14] L.-E. Cederman and L. Girardin, “Growing sovereignty: Modeling the shift from indirect to direct rule,” *International Studies Quarterly*, vol. 54, no. 1, pp. 27–48, 2010.
- [15] D. H. Bearce and E. O. Fisher, “Economic geography, trade, and war,” *Journal of Conflict Resolution*, vol. 46, no. 3, pp. 365–393, 2002.
- [16] S. J. Majeski, “Do exploitive agents benefit from asymmetric power in international politics?” *British Journal of Political Science*, vol. 35, no. 4, pp. 745–755, 2005.
- [17] M. J. North, N. T. Collier, J. Ozik et al., “Complex adaptive systems modeling with Repast Symphony,” *Complex Adaptive Systems Modeling*, vol. 1, no. 1, p. 3, 2013.
- [18] B. C. Schmidt, “Competing realist conceptions of power,” *Millennium: Journal of International Studies*, vol. 33, no. 3, pp. 523–549, 2004.
- [19] L. M. A. Bettencourt, “The origins of scaling in cities,” *Science*, vol. 340, no. 6139, pp. 1438–1441, 2013.
- [20] L.-E. Cederman, “Modeling the size of wars: from billiard balls to sandpiles,” *American Political Science Review*, vol. 97, no. 1, pp. 135–150, 2003.
- [21] I. A. Hatton, K. S. McCann, J. M. Fryxell et al., “The predator-prey power law: Biomass scaling across terrestrial and aquatic biomes,” *Science*, vol. 349, no. 6252, article no. 6284, 2015.
- [22] A. G. Haldane and V. Madouros, “The dog and the frisbee,” *Revista de Economia Institucional*, vol. 14, no. 27, pp. 13–56, 2012.
- [23] K. Zetter, *Countdown to Zero Day: Stuxnet and the launch of the world’s first digital weapon*, Crown, New York, NY, USA, 2014.
- [24] A. Greenberg, “Hacking team breach shows a global spying firm run amok,” in *Wired*, 2015, Hacking team breach shows a global spying firm run amok, <https://www.wired.com/2015/07/hacking-team-breach-shows-global-spying-firm-run-amok/>.
- [25] R. Axelrod and R. Illiev, “The strategic timing of cyber exploits,” in *American Political Science Association Annual Meeting*, p. 34, Chicago, IL, USA, 2013.
- [26] F. Bourguignon and T. Verdier, “Oligarchy, democracy, inequality and growth,” *Journal of Development Economics*, vol. 62, no. 2, pp. 285–313, 2000.
- [27] J. Sigholm, “Non-State Actors in Cyberspace Operations,” *Journal of Military Studies*, vol. 4, no. 1, pp. 1–37, 2013.
- [28] M. C. Libicki, “The nature of strategic instability in cyberspace,” *Brown Journal of World Affairs*, vol. 18, no. 1, pp. 71–79, 2011.

