7-2021

The Application of Complex Systems Science to Political Philosophy

Benjamin Nicholas House
University of Arkansas, Fayetteville

The Application of Complex Systems Science to Political Philosophy

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts in Philosophy

by

Benjamin Nicholas House
Wheaton College
Bachelor of Arts in English, 2015

July 2021
University of Arkansas

This thesis is approved for recommendation to the Graduate Council.

____
Warren Herold, Ph.D.
Thesis Director

Jacob Adler, Ph.D. Committee Member	Barry Ward, Ph.D. Committee Member

Abstract

Although complex systems science is relevant to problems of political philosophy, the intersection of these two disciplines has not been studied in depth. Complex systems are made up of multiple interdependent parts whose interactions create emergent properties. This interdependence makes these systems “fat-tailed”: low-probability events can have a major impact on the system. Complex systems engineers have formulated a series of rules of thumb for approximating an “evolutionary” environment. Contemporary human civilization is a complex system; because of this, governments need to become adaptable and approximate the evolutionary environment by fostering policy innovation while at the same time promoting mechanisms for altering or abolishing “toxic” policies. The best way to apply the techniques of complex systems engineering to government is for there to be a preference for smaller jurisdictions, decentralized governance, bottom-up policy creation, and discretionary policy implementation. However, the goal of making governments adaptable must be balanced against the other goals of government. Thus, there are situations in which larger jurisdictions, etc. are appropriate—primarily, cases which involve risk of grave moral harm or otherwise insoluble collective action problems. The complex systems science approach to political philosophy grounds many widely-held intuitions, but also provides some support for the political philosophy of Anglo-American conservatism.

Table of Contents

1. Introduction

Complex systems science has received increasing attention in the last few years, but despite its relevance to political problems, few attempts have been made to apply complex systems science to political philosophy. Many previous political philosophers lived in a time when the world was much less complex and interconnected. More recent political philosophers have experienced our interconnected, interdependent world—but have not systematically integrated the insights of complex systems science into political philosophy. My aim here is not to denigrate the valuable contributions by political philosophers of the past, but simply to point out a gap in the literature. The goal of this thesis is to systematically apply complex systems science to political philosophy, thus opening the door to further conversations between the disciplines. Thus, this thesis will necessarily be programmatic in its nature and broad in its scope.

I begin by sketching out what a complex system is, drawing upon a previously existing body of scientific literature—especially the writings of Yaneer Bar-Yam, the founder of the New England Complex Systems Institute. In short, complex systems are made up of multiple, interdependent parts which generate emergent properties that cannot be reduced to the properties of these parts. They are scale-variant—that is, their properties change with their scale. Larger complex systems are more efficient, but smaller systems are adaptable. Complex systems are fat-tailed; thus, outliers or "tail events" have a disproportionate impact on the entire system. This makes them vulnerable to "Black Swans." Naturally-occurring complex systems are created through evolutionary processes and exist in a dynamic evolutionary environment. Artificial complex systems, on the other hand, can be engineered in a similar way by applying a set of best practices.

I argue that contemporary human civilization is a complex system; thus, contemporary human governments are operating in a complex environment that is qualitatively unlike the environment of most of human history. In this environment, governments need to be adaptable; that is, they need to be able to generate more responses to the environment while filtering out potentially maladaptive responses. The best practices of complex systems engineering can be adapted to governments so that governments can be adaptable in a quasi-evolutionary manner—without taking on the ruthless, amoral, and unguided aspects of biological evolution. I argue that applying these practices leads to the following general rule: the burden of proof is on the one who proposes that jurisdiction should be larger rather than smaller, power more centralized rather than decentralized, policy making more top-down than bottom-up and policy implementation more uniform rather than discretionary. However, since adaptability is not the only goal of government, and since structural problems exist that can only be solved by increasing the scope of government, there are times when this general rule must be balanced against other considerations. Finally, I briefly sketch out how this approach is similar to the political philosophy of Anglo-American Conservatism, as it shares a preference for “small” government and a *trial-and-error approach* to policy.

2. What Are Complex Systems?

In ordinary terms, it may seem that a complex system is merely some sort of system that is complicated or difficult to understand, but in fact, the term *complex system* is a technical scientific term. Bar-Yam defines a complex system as “a system formed out of many components whose behavior is emergent, that is, the behavior of the system cannot be simply inferred from the behavior of its components. The amount of information necessary to describe the behavior of such a system is a measure of its complexity.” (*Dynamics of Complex Systems*

10). This definition is consistent with other definitions in the literature. Bar-Yam's collaborator Nassim Nicholas Taleb writes: "For our purposes, a complex system is one where, dynamically 1) interactions between parts can produce a different collective and individual outcome than when examined in isolation, 2) interactions are at least intermittently present" ("Scala Politica" 87). Jakobsson et al. write that "Complex systems are sets of interconnected elements whose collective behavior arises in a non-obvious way (and often counterintuitive and surprising way) from the properties of the individual elements and their interconnections". Finally, Mitchell defines a complex system as "a system in which large networks of components with no central control and simple rules of operation give rise to complex collective behavior, sophisticated information processing, and adaptation via learning or evolution" (13).

In the next portion of this thesis, I unpack and explain four key concepts in the study of complex systems: complexity, scale, interdependence, and emergence.

2.1 Complexity and Scale

What is complexity? Bar-Yam writes, "Loosely speaking, the complexity of a system is the amount of information needed in order to describe it. The complexity depends on the level of detail required in the description." (*Dynamics* 12). Suppose, in a moment of insanity, "we took a person and mixed up all the atoms so that they were no longer organized in any particular fashion." (Bar-Yam, *Making Things Work*, 56). We then poured these atoms into a vat—human soup. In this case, the movement of the atoms would be random and thus complex at the atomic level. It would take a large amount of detail to describe the movement of every single individual atom. At the finer scale of the atomic level, the vat is complex—it requires a lot of information to describe. However, if we looked at the vat of soup—or, if you prefer, a vat of purified water—it would appear to be simple, with little to no movement (to test this, pour yourself a glass of water and look at it). At a coarser[1] scale—the scale of ordinary human observation—the vat of liquid is simple. Although each atom is moving randomly, the overall mix would be at an equilibrium state (56-57). Thus, although it is complex at a fine scale (it would take a lot of information to describe the movement of each atom), it is simple at the coarser scale with which we would interact with it. The properties of a system at one scale are not always present at a different scale. Humans, for example, become more complex as we change the scale of our observation from that of a sociologist observing the movement of a person at a distance to the scale of an uber-scientist observing every single atom in a person's body (*Making* 55-56).

Let's imagine that we took the same atoms—the ones that make up a human being—and organized them so they were all moving the same direction, like an army marching in formation. Bar-Yam would call this a "coherent system" (57). This system is simple at both the atomic scale and at the coarser scale of everyday observation.

Finally, let's take a look at an actual person. Unlike the vat of human soup, or the formation of atoms, a human is complex at every scale. Because the movements of the human body and its constituent parts aren't totally random, the amount of complexity of the human at any scale never reaches that of the vat of soup at the atomic level; by the same token, because the movements of the human are not totally uniform (like the formation of atoms), the complexity of a human far exceeds that of a mass of atoms traveling in the same direction at the same speed (*Making* 55-58).

A "complexity profile," according to Bar-Yam and Siegenfeld, is "a plot of the system's complexity as a function of scale […] the scale of a behavior is equal to the number of coordinated components involved in the behavior" ("Introduction to Complex Systems Science"

---

[1] To avoid confusion, I will use the terms "fine" and "coarse" to refer to the scale of observation (e.g. observing the vat at the atomic level is observing it at a fine scale), and "large" and "small" to refer to scale in terms of size (e.g. A gallon jug of water is large scale, a pint glass of water is small scale).

3). Remember that the complexity of a system is different at different scales—the vat of soup or water is complex at the atomic scale, but simple at the scale of everyday interaction.

Bar-Yam writes, "To obtain the complexity profile, we observe the system at a particular length (or time) scale, ignoring all finer-scale details. Then we consider how much information is necessary to describe the observations on this scale" (*Dynamics* 14). Suppose we wanted to obtain a complexity profile for our vat at the atomic level. With the help of some powerful, though nonexistent, scientific equipment, we are able to observe and note the movement of every single atom in the vat. Then, we figure up how much information was necessary for this description. It is not necessary to calculate the amount of information it would take to log the movement of every single atom in order to realize that it would take a great deal of information to describe the movement of each atom. On the other hand, it would not take a great deal of information to describe the movement of the vat at the level of ordinary observation—unless it is disturbed, the water in the vat will just sit still.

There are tradeoffs between the scale of a system and its complexity. Bar-Yam writes,

> When parts are acting independently, the fine scale behavior is more complex. When they are working together, the fine scale complexity is much smaller, but the behavior is on a larger [coarser] scale. This means that complexity is always a trade-off, more complex at a large [coarse] scale means simpler at a fine scale. This trade-off is a basic conceptual tool that we need in order to understand complex systems (*Making Things Work* 58).

Again, we have seen this in the discussion of the vat, the formation, and the human. As coordination increases, complexity decreases. Suppose the atoms in the vat started to spontaneously form back into a human being—Swampman![2] The movement of the atoms would become less random, and thus less complex, at the finer scale. However, at the coarser scale, the human would be more complex than the vat of soup. If the movement of the atoms became more

[2] The example of Swampman is taken from Donald Davidson's paper "Knowing One's Own Mind," although Davidson uses the example for much different purposes in his paper.

and more uniform, to the point where Swampman started to morph into a formation of atoms all headed in the same direction, the complexity of the system would decrease at the higher levels. Bar-Yam writes,

> The intuition that complex systems require order is not unfounded: for there to be complexity at larger scales, there must be behaviors involving the coordination of many smaller-scale components. This coordination suppresses complexity at smaller scales because the behaviors of the smaller-scale components are now limited by the interdependencies between them. The tension between small-scale and large-scale complexity can be made precise: given a fixed set of components with a fixed set of potential individual behaviors, the area under the complexity profile will be constant, regardless of the interdependencies (or lack thereof) between the components. More precisely, the sum of a system's complexity at each scale (i.e., the area under its complexity profile) will equal the sum of each individual component's complexity. Thus, for any system, there is a fundamental tradeoff between the number of behaviors a system can have and the scale of those behaviors ("Introduction" 4).

Bar-Yam and Siegenfeld illustrate this with the example of a factory. A factory can produce a large number of the same kind of product (by coordinating the activity of the workers) or it can produce a small number of many different kinds of products (by allowing each worker to act relatively independently), but it cannot produce a large number of many different kind of products without making changes to the size of the factory (e.g. buying new equipment, hiring more workers, etc). ("Introduction" 4).

Another example: an "indie" game developer will develop, write, and program a game entirely by herself (perhaps with the help of pre-existing tools and assets). A small game development team of four people will divide up their responsibilities; perhaps our developer is now in charge of the artwork and graphics for the game. Although she does not have as much independence of action as before, she still has more independence compared to the developer at a large company such as Ubisoft or Bethesda, where a large team might work on the art and graphics for a game. The independent developer has more freedom of behaviors, but cannot produce a game on the scale of *Assassin's Creed: Valhalla* or *The Elder Scrolls V: Skyrim.* For comparison, the game *Undertale*, which was created by a "team" of one developer, takes up 200 megabytes of storage space ("Undertale"), while the newest (as of this writing) *Assassin's Creed* game, *Assassin's Creed: Valhalla*, takes up a whopping 50 gigabytes of storage space (256 times the size of *Undertale*) ("Assassin's Creed: Valhalla"). This impressive feat required coordination among the development team(s) at Ubisoft.

Although Ubisoft's larger team is able to create games of much greater scale, a smaller team or an individual developer has much more freedom in creating what they want to create. Bar-Yam and Siegenfeld point out that, "A corollary of the tradeoff between complexity and scale is the tradeoff between adaptability and efficiency" ("Introduction" 4). The factory that can create a large number of one product is efficient; the factory that can create a smaller number of many different products is adaptable. Bar-Yam and Siegenfeld continue:

> Adaptability arises when there are many possible actions happening in parallel that are mostly independent from one another, i.e., when the system has high complexity. Efficiency, on the other hand, arises when many parts of a system are all working in concert, so that the system can perform the task for which it was designed at the largest possible scale ("Introduction" 4).

The small development team has greater complexity; the workers act mostly independently of each other (assuming whomever is in charge is not a control freak) and have a large amount of freedom of action. Each member of the large development team at Ubisoft, on the other hand, has much less freedom of action; the team members must coordinate in order for Ubisoft to create and deliver their next 50 gigabytes worth of gaming to the market.

The relationship between complexity and scale has some interesting real-life consequences. As Bar-Yam writes,

> Due to the tradeoff between complexity and scale, a system with more adaptability will have a complexity profile with greater complexity but predominantly at smaller scales, while a system with more efficiency will have a complexity profile with lower complexity but extending to larger scales. Thus, a very efficient system will, due to its necessarily lower complexity, not be as adaptable to unforeseen variations within itself or its environment, while a very adaptable system, designed to handle all sorts of shocks, will necessarily have to sacrifice some larger-scale behaviors. ("Introduction" 4)

For example, if the team of four developers needs to make large-scale changes to their game, they can do so more quickly and with fewer costs than the team of 400 developers. To take the analogy further, a solo developer can do whatever she wants with the game at a relatively low cost, including scrapping the game entirely and moving to France to become a writer. The factory that makes a small amount of many different objects is less vulnerable to changes in the market; if the bottom drops out of the widget market, the adaptable factory will take a small hit, while the efficient factory, which only makes widgets, will suffer a large loss.

We shall see that this tradeoff has especially relevant consequences for political philosophy. In an increasingly complex world, our governments need to be "adaptable system[s], designed to handle all sorts of shocks" ("Introduction" 4). This point will be covered in more detail later in the thesis. For now, it is important to see why greater complexity, and the greater adaptability that comes with it, is desirable. Bar-Yam and Siegenfeld write,

> A determination of when complexity is desirable is provided by the Law of Requisite Variety: to be effective, a system must be at least as complex as the environmental behaviors to which it must differentially react. If a system must be able to provide a different response to each of 100 environmental possibilities and the system has only 10 possible actions, the system will not be effective. At the very least, the system would need 100 possible actions, one for each scenario it could encounter. […] Since complexity is defined only with respect to a particular scale, we can refine the Law of Requisite Variety: to be effective, a system must match (or exceed) the complexity of the environmental behaviors to which it must differentially react at all scales for which these behaviors occur. ("Introduction" 5)

Thus, an adaptable system is able to respond to a more complex environment. As the environment becomes more complex, the need to be adaptable increases.

An analogy can make this clearer. Suppose there are two musicians, David and Gianni. David's only significant income stream is live concerts. In fact, he sells out stadiums and makes millions of dollars from his live appearances. Gianni, on the other hand, has multiple income streams—he not only makes money from his live concerts, but also from album sales, streaming, merchandise, endorsements, production, guest appearances—even a little bit of acting. David is able to scale his business further than Gianni can, and make much more money than Gianni. However, if something happens to David's one income stream (say, a pandemic cancels all live concerts for the foreseeable future) he loses most of his income. Gianni, on the other hand, is less vulnerable to such shocks because he is more adaptable—if live concerts cease to be a source of income, he can focus on other sources. Similarly, because he has more skills—and thus a wider range of potential behaviors—he is able to deal with the different shocks of a complex environment. An event that would force David into unemployment—say, the loss of his ability to sing—would not be as devastating for Gianni, as he could use some of his other skills to make money.

Thus, we have seen complexity and scale. Let's now turn to the interesting features of complex systems—interdependence and emergence.

2.2 Interdependence and Emergence

Let's return to the stomach-turning example of the vat of human soup. The soup is made up of independent parts, namely, the atoms that once made up a human being. If we removed a selection of the atoms—say, one cup of human soup—there is no change in the properties of the soup as a whole. The vat remains in equilibrium.

Similarly, if we removed a cup's worth of atoms from the formation of atoms moving in a straight line, the nature of the formation would not change. The rest of the atoms would keep moving in a straight line. The components of this coherent system are coordinated, but they are not interdependent.

Now suppose that we removed a cup's worth of material from a person—perhaps scooping out their prefrontal cortex from their brain. This would cause major changes in the person's brain, and in the person as a whole. The parts of the brain are interdependent—they rely on each other to function properly. Removing any part of the brain will impair the function of the brain at best, and destroy the function of the brain at worst. The same thing holds for the human body as a whole—removing an arm or a leg will change the body much more than removing a cup of soup from the vat will.

Describing an "attractor network" (which is a model similar to the neural networks of the brain), Bar Yam writes,

> In the model, there are synapses between each neuron and every other neuron. If we remove a small part of the network and look at its properties, then the number of synapses that a neuron is left with in this small part is only a small fraction of the number of synapses it started with. If there are more than a few patterns stored, then when we cut out the small part of the network it loses the ability to remember any of the patterns, even the part which would be represented by the neurons contained in this part (*Dynamics* 11).

The individual components of the network are interdependent—the network cannot function properly if any of them are missing.

One more example of interdependence comes from software, especially video games. Every computer program you use is made up of many lines of code. Each part of the code that makes up a program is interdependent. Cutting out a part of the code could cause the program to malfunction or crash. Similarly, adding to the code can do so as well. Some computer games, such as *Doom* or *The Elder Scrolls V: Skyrim* are “moddable”—hobbyists can change and add to the game, whether by modifying the source code directly or by using a creation tool provided by the original developers. Longtime players of these kinds of games will recall how poorly-made or incompatible mods can cause a game to malfunction or crash.

Interdependence allows for emergence. The evocative phrase “the whole is more than the sum of its parts” roughly captures the idea of emergence. More specifically, emergence occurs when property P in system S is not present in any one of the parts that makes up S, but “P emerges from interactions of component parts of S.” (Norman, “Part I”). For example, Norman points out that neurons in the brain (a network similar to an attractor network) have two states, “on” and “off,” without the ability to recognize patterns. Neurons do not have a “pattern-recognition property” (“Part I). “However”, he writes, “when these neurons are connected into networks, they are able to recognize complex patterns, such as human faces” (“Part I”). Pattern recognition in the brain is emergent—it emerges from the interaction of the different components of the brain, but is not a property of any one part of the brain. Similarly, the body has emergent properties that are not properties of any one part of the body. One can say, for example, that John is healthy, but “health” is not a property of any one part of his body, but an emergent property of the entire body created by the relationships between the different parts of the body.

Computer programs also have emergent properties that emerge from the interaction of their parts. No program can run without its code, but the experience of sending emails on Gmail, writing this document in Microsoft Word, or playing a game like *Doom* or *Skyrim* game cannot be reduced to the properties of the individual lines of code. Each program is actualized by the interaction of the different pieces of code with each other, and by the interaction of the program with the hardware it is being run on. All these pieces are interdependent. Without the hardware to run the program on, for example, it is impossible to use the program. Similarly, if an important piece of the code is missing or is corrupted, the program will crash or be unusable.

It is important to note that just "because the collective behavior is not readily understood from the behavior of the parts" does not imply that "that the behavior [of the system] is not captured by the behavior of the parts" (*Dynamics* 10). Rather, "The collective behavior is […] contained in the behavior of the parts if they are studied in the context in which they are found" (*Dynamics* 10). This is the case because the emergent behavior emerges not from the parts themselves, but from the relations between the parts. If we look at the parts in isolation, we cannot understand the collective behavior from them; but we can understand (at least to some extent) how the collective behavior arises from the parts as they interact with one another. If this were not the case, then developing a computer program would be impossible—the developers need to have a rough idea of how the different parts of the code will interact to form the program. Similarly, an understanding of the body would not be possible if we could not understand collective behavior in terms of the behavior of interacting parts.

However, the fact that emergent behavior cannot be understood by observing the behavior of the parts by themselves (but of the parts in their context in a system) seems to imply that our method of understanding complex systems should tend in an empiricist, trial-and-error direction rather than a rationalist one. As I shall argue later in the thesis, this is broadly correct—there are reasons for believing that the behavior of complex systems cannot be predicted.

One final note here: it is true but trivial that non-complex systems, such as the vat and the formation mentioned earlier, also display emergence. This kind of emergence is what Bar-Yam calls "local emergence […] where collective behavior appears in a small part of the system"

(*Dynamics* 10). An example of local emergence would be the properties of gasses. Bar-Yam writes: "Two emergent properties of a gas are its pressure and temperature. The reason they are emergent is that they do not naturally arise out of the description of an individual particle" (*Dynamics* 10). These properties are "local emergent properties," however, because "We can take a very small sample of the gas away from the rest and still define and measure the (same) pressure and temperature" (*Dynamics* 10). The individual parts of the gas are not interdependent; thus, the emergence is only local. "Global emergence," occurs "where collective behavior pertains to the system as a whole," and is the kind of emergence that occurs in a system made up of interdependent parts, such as computer programs, the human body, and, as we shall see, human civilization (*Dynamics* 10).

Thus far I have explained what a complex system is, and elucidated the terms *complexity*, *scale*, *interdependence,* and *emergence.* I turn now to the nature of complex systems.

3. How Complex Systems Work.

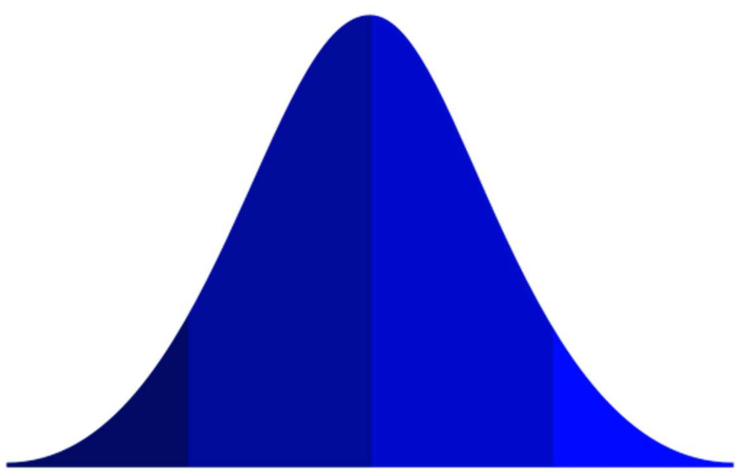

Figure 1: Graph of a normal distribution

This graph of a probability function illustrating normal distribution, also known as the Bell Curve, looks like a large, evenly spaced hump. The very middle of the bell curve—the peak of the hump--sits over the mean of probabilities. On either side, as the curve tapers off, we can see that the parts of the distribution far from the mean look like tails—hence, the term "tails" is used to describe the parts of a probability distribution that are far from the mean. "Tail risk" is generally understood in finance as the risk of a "move of at least three standard deviations […] from the mean" (Akoundi and Haugh 1). In non-financial, non-quantifiable domains, one can understand tail risk to mean the risk of a (negative) low-probability event, also known as a "tail event."

Tails can be either "thin" or "fat." To understand what this means, it will help to introduce two concepts invented by Nassim Nicholas Taleb—"Mediocristan" and "Extremistan." "In Mediocristan," Taleb writes, "when a sample under consideration gets large, no single observation can really modify the statistical properties" ("Statistical Consequences of Fat Tails 22). Human strength, for example, is a domain in Mediocristan. The average untrained adult male has a one rep maximum deadlift of 155 pounds. (Duquette and Walker-Ng) I, on the other hand, have a one-rep max of 310 pounds. Let's calculate the average strength of a group of 99 average men plus myself. With my 310-pound lift added, the average increases from 155 to roughly 156 pounds. Suppose now I am added to a group of 999 lifters. My contribution now changes the average deadlift from 155 pounds to 155.155 pounds. Finally, suppose I am the one man out of a million who can deadlift 310 pounds—a god in human form, surrounded by puny insects who struggle to lift 155 pounds. The average deadlift of the entire group is 155.000155. My "contribution" to the average becomes smaller and less significant as the size of the group increases. Even if I was replaced by Hafthor Bjornsson, the current holder of the world record for deadlifting, his one-rep max of 1,104 pounds would only increase the average to 155.000949 pounds. Thus, Mediocristan is *thin-tailed* because the low-probability events (such as the person who can lift over a thousand pounds) do not affect the properties of the sample.

On the other hand, Taleb writes that, “In Extremistan, the tails (the rare events) play a disproportionately large role in determining the properties [of the sample]” (*Statistical Consequences* 22). Thus, probability distributions in Extremistan are *fat-tailed*—in other words, low probability-events (“tail events”) can affect the properties of the sample. Somewhat counterintuitively, “The fattest tail distribution has just one very large extreme deviation, rather than many departures” from the median (*Statistical Consequences* 23). This is counterintuitive because this tail would look long and thin if the probability function were graphed, yet is referred to as a “fat tail.”

Suppose instead of calculating strength, we calculated wealth. Let's assume for the sake of argument that the average salary for all humans is $18,000 per year. This is not entirely accurate, but it will work for our purposes. Suppose I make $35,000 a year, as was once the case; if we add me to the pool of 99 average wage-earners, the average will increase to 18, 170. However, suppose instead of me, we add Bill Gates to the pool, and suppose Bill Gates makes $3,000,000,000 per year (again, the figure is not entirely accurate). Gates raises the average income of the group to $30, 017,820. If he is added to the group of 999 average wage-earners, he raises the average to $3,017,982; if added to the group of 999,999 wage-earners, he increases the average salary to roughly $21,000. Wealth, unlike human strength, exists in Extremistan. Extremistan is fat-tailed; the high probability events affect the properties of the sample. Thus, even if there exists only one super-wealthy person in the sample, wealth still falls into the domain of Extremistan, and is thus fat-tailed.

Some more examples will be helpful. Suppose I invest $30 and the market follows a normal distribution with a standard deviation of $5. Most likely I will end up with somewhere between $20 and $40. If the market follows a normal distribution, then I don't have to worry about losing a lot of money—nor will I gain a large amount either. Most of the gains or losses will be within one standard deviation of the mean. In a fat-tailed distribution, on the other hand, a tail event could result in a huge gain—or a huge loss. If an event that is "in" the right side of the tail occurs, I could make a lot more money from my $30 investment; on the other hand, it is quite possible that I could lose all my money. Thus, I need to take some measures to protect against negative tail events (for example, setting a stop-loss order to prevent myself from losing money, or diversifying my portfolio).

Complex systems exist in the domain of Extremistan because they are fat-tailed; in these systems, "low-probability" events have a greater impact. The fat-tailed nature of complex systems derives from the interdependent nature of their parts. As Bar-Yam and Siegenfeld write,

> When the components of a system are independent from one another above a certain scale, then at much larger scales, the magnitudes of the fluctuations of the system follow a normal distribution (bell curve), for which the mean and standard deviation are well defined and for which events many standard deviations above the mean are astronomically improbable. Interdependencies, however, can lead to a distribution of fluctuations in which the probability of an extreme event, while still small, is not astronomically so. Such distributions are characterized as fat-tailed. ("Introduction" 8-9).

Not only are these extreme events more probable, but they affect the properties of the entire system (think of the impact of Bill Gates on the modern world vs. the impact of Hafthor Bjornsson). This leads to significant consequences.

3.1 Black Swans

The fat-tailed nature of complex systems means that they are more vulnerable to the kinds of events that Taleb calls Black Swans and Gray Swans. According to Taleb, a Black Swan is an event with three characteristics:

> First, it is an *outlier,* as it lies outside the realm of regular experience, because nothing in the past can convincingly point to its possibility. Second, it carries an extreme impact (unlike the bird). Third, in spite of its outlier status, human nature makes us concoct

> explanations for its occurrence *after* the fact, making it [seem] explainable and predictable (*Black Swan* xxii).

Similarly, Gray Swans are "rare and consequential, but somewhat predictable" events (*Black Swan* 37). As Taleb writes, "If you know that the stock market *can* crash, as it did in 1987, then such an event is not a Black Swan"—although it is still a tail event with outsized impact. The terrorist attacks on September 11th are examples of Black Swans. Previous hijacking attempts had not involved *crashing* a plane into a building, Thus, the attacks were *outliers;* the possibility of this specific kind of an attack was not recognized even by the federal agencies that dealt with aircraft hijackings (National Commission 17). The attacks carried with them an "extreme impact," resulting in many political changes. Finally, while the attacks were unpredictable, it was relatively easy to trace the chain of events after the fact—making them *seem* predictable.

The consequences of climate change, on the other hand, are an example of a Gray Swan. We cannot predict the consequences in an absolutely precise way, but we know enough to be able to predict a range of possible consequences. One can come up with a contingency plan for a Gray Swan, but it is impossible to plan for a Black Swan.

Black Swans and Gray Swans are unpredictable and difficult to predict respectively due to the nature of complex systems. Taleb illustrates this by citing an argument from Henri Poincare:

> [A]s you project into the future you may need an increasing amount of precision about the dynamics of the process you are modeling, since your error rate grows very rapidly. The problem is that near precision is not possible since the degradation of your forecast compounds abruptly—you could eventually need to figure out the past with infinite precision. Poincare showed this isn't a very simple case, famously known as the "three body problem." If you have only two planets in a solar-style system, with nothing else affecting their course, you may be able to indefinitely predict the behavior of these planets, no sweat. But add a third body, say, a comet, ever so small, between the planets.

> Initially the third body will cause no drift, no impact; later, with time, its effects on the two other bodies may become explosive. Small differences in where this tiny body is located will eventually dictate the future of the behemoth planets (177).

Remember that the interaction of the components in complex systems leads to non-linear results. This makes precise predictions about complex systems impossible for humans, because predicting the motion of a complex system requires knowing more and more variables as the complexity of the system increases. Taleb gives an example taken from the work of Michael Berry: predicting the movement of balls on a pool table requires and amount and precision of knowledge that increases non-linearly as the successive movement of the balls take place—"to compute the fifty-sixth impact, every single elementary particle of the universe needs to be present in your assumptions!" (*Black Swan* 178). This means that precisely predicting the behavior of complex systems with more variables than pool table (such as the weather, or the economy) is not possible. As Taleb writes, "In a dynamical system, where you are considering more than a ball on its own, where trajectories in a way depend on one another, the ability to project into the future is not just reduced, but is subjected to a fundamental limitation" (*Black Swan* 178). Thus, we cannot take a rationalist, *a priori* approach to understanding complex systems; an empiricist, *a posteriori,* trial-and-error method is the only way to come close to understanding a complex system.

Note that the intractability of complex systems does not mean that we cannot know anything about them. Our inability to make precise predictions about the behavior of complex systems does not imply the inability make any predictions at all. One can, for example, predict that another terrorist will occur sometime in the future, based on the fact that these have occurred in the past (although this is likely to be true but trivial). The inability to predict also precisely does not imply that one must be helpless in the face of unpredictability—one must simply adjust one's strategy, moving from preparing for specific events to identifying and remedying systemic weaknesses. For example, while you cannot predict what health problems you will face in the future—and thus, cannot prepare for a specific ailment—you can improve your overall health, and thus minimize (though never eliminate) your risk of high-impact health problems. Similarly, the inability to predict individual terrorist attacks does not mean that nothing can be done to prevent or respond to terrorist attacks. If (some of) the security measures implemented after 9/11 had not been implemented, other terrorist attacks of a similar kind would have surely occurred.

3.2 System Failure

The interdependent nature of complex systems and their propensity to fat-tailed behavior leads to the possibility of (undetected) system failure. Bar-Yam and Siegenfeld write:

> One danger of interdependencies is that they may make systems appear more stable in the short term by reducing the extent of small-scale fluctuations, while actually increasing the probability of catastrophic failure. This danger is compounded by the fact that when underlying probability distributions have fat tails (a situation made more likely by interdependencies), standard statistical methods often break down, leading to potentially severe underestimates of the probabilities of extreme events ("Introduction" 9)

For example, suppose all the appliances in your house are "smart" appliances, hooked up to the Internet of Things and to each other. You can now set the lights to dim and the temperature to lower whenever it's time to go to bed! The hassle of programming each device separately has been removed—you can access everything from an app on your phone. Until, of course, your Internet Service Provider cuts off service because you were illegally downloading philosophy papers. If the internet fails, all your "smart" devices are useless; a person who has a series of independent devices that did not rely on the internet would be in a much better situation.

Similarly, the recent widespread power outages in Texas have shown the danger of dependency on electrical power ("Burst Pipes"). Electricity reduces the extent of small-scale fluctuations by providing devices that work more reliably and smoothly on a regular basis than their non-electrical counterparts. It is much easier to keep one's house at a steady temperature by setting the thermostat than by using a wood stove in the winter and opening the windows in the summer. However, when multiple components of a system are reliant on electrical power, the failure of the electrical grid can result in catastrophic failure for the system.

Bar-Yam and Siegenfeld give another example:

> As a thought experiment, imagine 100 ladders, each with a 1/10 probability of falling. If the ladders are independent from one another, the probability that all of them fall is astronomically low (literally so: there is about a times higher chance of randomly selecting a particular atom out of all of the atoms in the known universe). If we tie all the ladders together, we will have made them safer, in the sense that the probability of any individual ladder falling will be much smaller, but we will have also created a nonnegligible chance that all of the ladders might fall down together (9)

The interconnected nature of our global society leads to more situations of this type. Bar-Yam and Siegenfeld point out that market crashes (like the 2008 financial crisis) and global pandemics (like COVID-19) are examples of these kinds of systemic breakdowns. However, there are others. In the sci-fi series *Dune* by Frank Herbert, the entire galactic empire relies on a substance called "the spice," a mind-altering chemical which, among other things, powers the "Navigators" who guide spaceships through the galaxy. The spice is only produced on one planet, Arrakis. Thus, whoever controls the planet Arrakis has outsized political influence on the entire galaxy (a major plot point in the books).

Investor and writer Ben Hunt noticed a real-life parallel to the situation in the Dune novels when computer company Intel ordered a large number of chips from the Taiwan-based TSMC ("Taiwan is Now Arrakis"). He wrote,

> The world's principal supplier of semiconductors – the spice of OUR global empire – is now Taiwan. […]There is no future where the United States can both maintain its

> existential national interests and allow the world's principal supplier of semiconductors to come under the direct political control of China.
>
> And there is no future where China can both maintain its existential national interests and allow the world's principal supplier of semiconductors to remain outside its direct political control ("Taiwan Is Now Arrakis").

Suppose that Taiwan is the only supplier of semiconductors. Any major event that happens to Taiwan, such as political instability or a natural disaster, will affect the supply of semiconductors—and thus will affect every industry that depends on computers and the semiconductors used to build them. If Taiwan was to end the exporting of semiconductors, it could cause the collapse of entire industries.

Similarly, the Apollo Masters Corporation was the only producer of lacquer discs in America, and one of only two in the world. The fire that occurred at their factory has serious consequences for the vinyl record industry. For example, vinyl pressing plants had to place new limits on which and how many records they could press; some moved to a more expensive pressing technique known as "direct metal mastering" (Grant).

On the other hand, a lack of interdependencies can result in greater resilience to tail events. For example, a farmer who grows her own food will not be affected by problems with supply chains the way someone who gets their food from a grocery store will be. However, she will also have to put out more work and will not have access to the same amount or variety of food that a grocery shopper will. The takeaway is that levels of interdependency require tradeoffs—one must be aware of the risks as well as the benefits of increasing complexity.

The fat-tailed nature of complex systems leads to the possibility of rare events (Black or Gray Swans) with *cascading effects.* Suppose I break my arm. In this case, the event only negatively affects me; perhaps others are mildly affected by it; some, such as the doctor, may mildly benefit from it. If I die, others experience negative effects, but those effects becomes weaker and weaker as they spread further from me—those closest to me are affected less than those further away.

On the other hand, if I am the first person to catch a new and highly contagious disease, the effects of my catching the disease and spreading it will become stronger as they continue, like a series of falling dominoes where each domino becomes larger. Similarly, the effects of climate change can cascade: increasing temperatures create environmental effects that result in reduced food sources, which can in turn lead to political turmoil and war.

The proper response to the possibility of these kinds of events is to get rid of the systemic dynamics that create the conditions for these events occurring. Bar-Yam and Siegenfeld write,

> When such crises do occur, they are often attributed to proximate causes or chains of events, and measures are then implemented to ensure that those particular chains of events will not occur again. But unless the underlying systemic instabilities are addressed, another crisis is bound to happen sooner or later, even if its precise form cannot be predicted. (9)

Thus, it is important to pay attention to the strengths or weaknesses of the system as a whole rather than to attempt to prevent future events that are similar to the past event. For example, a naïve approach to dealing with terrorist threats would be trying to prevent the same kind of attack from happening—for example, focusing on preventing another event where planes are flown into a skyscraper. An example of a systemic approach would be focusing on de-radicalization of potential terrorists—the goal would be to change the system to reduce the probability of future terrorist attacks of any kind rather than one specific kind of terrorist attack. It is not likely that the same exact kind of terrorist attack will happen again, which is why focusing on system-wide vulnerabilities is a necessity.

3.3 Evolutionary Processes

The complex systems seen in nature, such as cells, organisms, and ecosystems, were not assembled all at once according to a blueprint, but instead emerged through evolutionary processes of trial and error. Bar-Yam writes, “The formation of complex systems, and the structural/functional change of such systems, occurs through a process of adaptation, especially through evolution” (“Strategy” 9). For example, the complex system that is the human body developed through the process of evolution, “the adaptation of populations through intergenerational changes in the composition of the population.” (“Strategy” 9). The human organism is the way it is today because its ancestors had adaptations that allowed them to reproduce. Organisms and species that failed to adapt to their environments exited the gene pool.

In addition to the multigenerational process of evolution, complex systems can adapt during their lifetime through the process of learning, which Bar-Yam defines as a “process of adaptation of a system through changes in its internal patterns, including, but not exclusively, the changes in its component parts” (ibid 9). Think, for example, of the way the brain learns. As new behaviors or ideas are learned, the arrangement of the synapses in the brain changes. Neural connections are re-wired and strengthened or weakened in response to environmental changes.

However, not all complex systems are “engineered” by nature in such a way. Some complex systems are created by humans—for example, the market, human cultures, software, etc. These complex systems are not the result of an evolutionary process the way that an organism or an ecosystem is. They can be shaped and molded by human actions—and these actions can be more or less successful.

Observing evolutionary processes can help us formulate a set of “best practices” or “practical principles (Norman, “Part III”) for engineering and working with complex systems.

Nature's processes have been time-tested over billions of years and have resulted in systems far more complex than humans have built—the brain, for example, is more complex than any computer. So we can take a few cues from Nature and experience in order to formulate some guidelines for working with complex systems. These principles are not a systematic blueprint for building a complex system—rather, they are suggestions for creating an "evolutionary environment in which problem-solving systems can evolve." (Norman, "Part III").

In an "evolutionary process," "successful changes are copied and unsuccessful changes are not" ("Introduction" 10). Recall the system of the vat: if we dump a rock into the vat, it will change, but will quickly return to the equilibrium state. Systems embodied in an evolutionary process, however, can experience qualitative changes. For example, an animal that possesses a "superior" adaptation will spread its genes; an animal with a severe maladaptation will not. When we learn useful new information, we keep a "copy" of that information in our brain, so to speak; when we debunk false information, we stop "copying" that information. Another example: in a system where restaurants compete for customers, restaurants with bad service will generally be filtered out of the system; as restaurants compete, those with innovations that satisfy customers' needs (real or imagined) will be more likely to succeed. Imagine a world where restaurants with poor service were continually bailed out by the government. No restaurant would have an incentive to offer a better product than any other restaurant.

Note, however, that the evolutionary filtering process is not 100% reliable. No animal is 100% optimized to its environment. Similarly, even in a relatively free market system, some companies that offer good products fail and some companies that offer shoddy products succeed. Evolutionary processes filter out some (a significant amount) of unfit units and provide a system where units become (somewhat) more optimized.

Bar-Yam notes that systems that are "[embodied] in some sort of evolutionary process […] in which successful changes are copied and unsuccessful changes are not" will actually be strengthened and "benefit from uncertainty and variability" ("Introduction" 10). In other words, these evolutionary systems are "antifragile"—they gain some benefits from disorder (Taleb *Antifragile* 12-13). Taleb gives an exhaustive list of the "Extended Disorder Family (or Cluster)" in *Antifragile;* it consists of:

> (i) uncertainty, (ii) variability, (iii) imperfect, incomplete knowledge, (iv) chance, (v) chaos, (i) volatility, (vii) disorder, (viii) entropy), (ix) time, (x) the unknown, (xi) randomness, (xii) turmoil, (xiii) stressor[s], (xiv) error[s], (xv) dispersion of outcomes, (xvi) unknowledge (*Antifragile* 13).

Evolutionary systems will benefit from (some) disorder because disorder can both create successful changes and weed out unsuccessful ones. For example, a rock will generally retain the same properties over time-any shocks or stressors will weaken or break it. A human, on the other hand, can gain in strength by lifting the rock, building her muscles by exposing them to stress—but not too much stress. Thus, humans (or at least their musculo-skeletal system) are "antifragile."

Unlike the fragile, which is broken by stress, the antifragile becomes stronger due to stress. In the body's case, slight stressors and variability, such as those experiences in exercise, will strengthen the body. Exposure to randomness and complex information can strengthen the mind as well—in the paper "Environmental Complexity: Information For Human Environment Well-Being," Bar-Yam and Alice Ware Davidson studied the living environments of elderly people in retirement communities. They discovered that "more complex environments correlated with the higher cognitive function and more robust locomotor activity of those living in the community" (1). Complex environments contained more uncertainty and unpredictability than the simple environments, which "exercised" the brain in a manner analogous to exercising a muscle.

Similarly, the market benefits from some measure of uncertainty and disorder. Bar-Yam writes,

> Competitive market economies provide another example of how systems can thrive on uncertainty. Due to our ignorance of which will succeed, many potential innovations and businesses must be created and improved upon in parallel, the successful ones expanding and the unsuccessful ones failing. The successful among these can then be improved upon in the same manner—with many approaches being applied at once—and so on ("Introduction" 10).

On the other hand, too much disorder will break the system down, changing the system from its complex, dynamic state back to a state of equilibrium, like the vat of human soup we looked at earlier. Complex systems such as the human body have an upper bound of how much disorder they can take. One can gain a tolerance to lifting heavy weights by lifting progressively heavier and heavier objects—but one cannot gain a tolerance to puncture wounds by shooting oneself with increasingly higher-caliber bullets, because the disorder introduced by being shot is too great for the system to handle. Similarly, dinosaurs were able to survive 165 million years filled with changes and fluctuations, but were wiped out when a meteor quickly introduced a change so large they were unable to adapt to it. If the climate changes introduced by the meteor had taken place gradually, dinosaurs as a whole would have been able to evolve and adapt. Similarly, a society, that faces a large enough stressor—in other words, that faces too much disorder—will collapse. Thus, human-engineered complex systems need to face enough disorder to successfully "evolve" and improve, but not so much that they disintegrate.

3.4 Best Practices for Engineering Complex Systems

Based on the nature of complex systems, complex systems scientists have formulated a series of “best practices” or “practical principles” (Norman, “Part III”) for engineering complex systems. These practices are not inviolable scientific laws or mathematical theorems, but rather helpful rules of thumb. The goal of these practices is to approximate the evolutionary environment of naturally occurring complex systems. There is no “canonical” list of these best practices; I will list five, but this list is not exhaustive.

The first best practice is to “Foster (non-toxic) variety” at a small scale (Norman, “III”). A system with a limited amount of behaviors will be unable to react to the environment—however, a system with a large and expandable amount of behaviors will be able to react and adapt to the environment (provided that it doesn’t perform a behavior that would destroy it.) The reason why this variety must be fostered at a small scale is due to the complexity-scale tradeoff discussed earlier (Bar-Yam, “Introduction” 11). It is not possible to create such variety at a large scale.

One example of this best practice is Google’s “20% Time” policy. At one point in time, Google had an informal policy of allowing employees to spend 20% of their time working on side projects (Tate). This policy resulted in the creation of, among other things, Gmail and Adsense, which have both brought the company an enormous amount of revenue (ibid). The 20% policy has also been emulated by a number of other companies, including Facebook and LinkedIn (ibid). While it would be inefficient for the entire company to explore a large variety of projects, it is possible to allow employees to explore these projects, thus exploring the space of possibilities and giving the company more options.

The second best practice is “Expose to the ecological early” (“Part III”). Potential choices need to be tested in the real world; choices that seem successful in theory may prove to be quite unsuccessful in practice. Norman notes that “Exposure to the real problem environment the system is supposed to operate in during development/evolution of varieties will buffer against over-designing, and provide an opportunity for the maturation of stems that can handle the true complexity of their task” (“Part III”). For example, large companies such as McDonalds often test new products in a smaller market before rolling them out to the entire franchise (“How do you product test new products”). In the event that the product is unsuccessful, the company only suffers a small loss. If McDonald’s rolled out a new product to the entire market before testing it, they could lose a lot of money. Because of the fundamental limitations in understanding complex systems, the best filtering mechanism for potential choices is exposing them to disorder early.

The third best practice is “Detect and fail fast, and local, the toxic” (Norman, “Part III”). Just as successful choices should be replicated, unsuccessful choices should be jettisoned as quickly as possible. The reason this process should be fast and local is to keep “harmful varieties” from “becoming systemic.” (Norman “Part III”). Recall the possibility of cascading effects—we want to avoid a negative effect from cascading and damaging or even overwhelming the system.

The fourth “best practice” is to “allow for a means of communication between various parts of the system so that successful choices are adopted elsewhere and built upon.” (Bar-Yam “Introduction” 12) Without his kind of communication, successful choices will remain isolated and fail to replicate. For example, if Google employees were unable to communicate with each other, they would be unable to refine the projects that they used their 20% of free time on. On a larger scale, if Google was entirely opaque, other companies would not be able to adapt and build upon their innovations.

The fifth best practice is “Resist the temptation to scale quickly a promising solution,” noting that a short time-scale will hide potential “malignancies,” while at the same time successful solutions can only prove themselves overtime (Norman, “Part III”). To return to the previous example, if the costs of rolling out a new product to the entire market are high, then McDonalds will want to test the product for some time to make sure that the product is cost-effective. Initial successes may be flukes due to many factors. Similarly, the initial positive effects of a new medicine may be due to the placebo effect or other factors. Multiple tests in multiple scenarios are necessary to determine whether the medicine is effective.

The upshot of these five best practices is that a small-scale, trial-and-error approach is the best way to approximate naturally occurring evolutionary processes, and to gain the benefits of such processes. The preference for trial-and-error may seem true but trivial, but it is not. This approach is adopted, not due to personal preference or unwillingness to pursue other options, but due to the nature of the scenario. As we have seen, the behavior of complex systems cannot be accurately predicted, so an *a priori* approach to engineering complex systems will not work. A trial-and-error approach is appropriate to complex systems engineering in a way that it is not be appropriate in the case of a math problem, where the best way to find the solution is to apply a set of rules. This is because the purpose of the trial-and-error approach is to approximate the naturally occurring evolutionary processes which are responsible for the many intricate and stable arrangements we see in nature. Thus, the trial-and-error approach is not merely a piece of folksy wisdom, but an attempt to replicate the most successful existing complex systems.

4. The Application of Complex Systems Science to Political Philosophy.

4.1 Human Civilization is a Complex System

Complex systems science is relevant to political philosophy because contemporary human civilization is a complex system. Therefore, the best practices from complex systems engineering can be adapted to the context of government in order to ensure that our societies are adaptable and resilient.

It may seem trivial to say that human civilization is more complex than it has been in the past. Nevertheless, there are good reasons to believe that human civilization in 2021 is a complex system in the way that we have defined, and that this makes our contemporary situation qualitatively different than the state of the world in, say, 1521. Recall that a complex system is a system made up of interdependent parts whose interactions create emergent properties that are not present in the individual parts. The more parts and interactions, the more complex the system is. The properties of these systems are scale-variant—they change as the system (number of interacting parts) becomes smaller or larger. Complex systems are fat-tailed; rare, low-probability events can have an outsized impact on the entire system. Thus, complex systems are vulnerable to tail risk.

These features of complex systems are features of our world today but were not features of human civilization in the past. While individual societies are complex systems, human civilization throughout most of history was not a complex system because the different component parts of human civilization were not strongly interdependent. A member of a tribe in Stone Age Papua New Guinea was interdependent with the other tribe members—his death would affect the properties of the tribe as a whole. However, the fate of this man, or even his entire tribe, was not relevant to a tribe in New Mexico. Someone viewing human civilization during the Stone Age at the coarse scale of the entire earth (say, an alien viewing it from a satellite) would see it as a simple system, made up of multiple parts that do not interact and exist in an equilibrium state. Even when societies became more interdependent due to advances in transportation and communication, their interactions were limited in amount and frequency by the slow nature of communication and transportation. Because societies were relatively independent, the kinds of interactions that lead to emergence did not regularly obtain. Thus, human civilization in the past was not a complex system.

However, as the speed and extent of transportation and communication increased, the scope of markets and inter-government relations increased as a consequence. This led to nations and the citizens of those nations being strongly interdependent in a way that they were not in previous times. Human civilization in the post-Industrial era has a significantly greater complexity profile than human civilization during most of human history. That is to say, contemporary human civilization requires much more information to exhaustively describe, because it has far more interdependent, interacting components. While an observer of Stone Age or Medieval earth would view a simple system of isolated parts, similar to the atoms in a vat of water, an observer of contemporary earth would view a complex system of interacting parts, similar to the human body or an ecosystem

The complex nature of contemporary human civilization means that this system is fat-tailed, with rare events having an outsized impact on the entire system. Here are a few examples demonstrating this.

Supply chains: Throughout most of human history, nations have existed in a state of relative economic independence. In today's global economy, countries are interdependent because they provide essential or economically significant resources for each other. One country may export a resource to another country and in turn import another resource from that country. I have presented earlier examples of interdependence in the vinyl and superconductor industries—the removal of a single factory or country can cause vast disturbances in the supply of a product, which in turn can drive political and cultural changes. A similar example comes from the recent blockage of the Suez Canal, in which an accident involving one ship created "cascading effects" on the global economy (Flynn, qtd. in Stevens, "Suez Canal"), including an increase in oil prices ("Egypt's Suez Canal blocked"). In earlier stages of human civilization, the blockage of one canal would have had primarily local effects, because societies were more independent. Similarly, the lack of advanced technology meant that it was not necessary for nations to import goods on a large scale—one does not need a wide-ranging supply chain to build a chariot. Today, the interdependent nature of the global economy leads to the ability to manufacture and distribute advanced technology on a large scale, but the dependence on advanced technology makes the system vulnerable to the kind of bottlenecks described above.

Political Instability: In earlier times, lack of swift communication and transportation limited the extent of political alliances, and thus limited the interdependence of nations. Assassinations, insurrections, and succession crises occurred, but their impact was spatially limited. The assassination of a "Big Man" in the society of Stone Age New Guinea would have few effects beyond the tribe and its neighbors. The assassination of a nobleman in Renaissance Europe would have a greater effect but would still be primarily a local phenomenon. News of this event could take weeks or months to reach the rest of Europe. In contrast, the assassination of Archduke Franz Ferdinand was the catalyst for the First World War, which was a global conflict. The reason for this is that the countries of Europe had become so interdependent that events in Serbia could affect the rest of the continent. This interdependence is not possible without the technology that allows swift communication and transportation, thus making more extensive political alliances and economic interchange possible. Additionally, the increase in the destructive power of weapons has changed the face of warfare and thus of international relations. A military's destructive capacity is no longer limited by its manpower, allowing smaller nations like North Korea, Israel, and Iran to have a greater influence due to the possession of advanced military technologies. Thus, the domain of international relations is fat-tailed; an individual act of aggression by a small country can have an outsized impact on the entire system.

Pandemics: Due to the slow and intermittent transportation in previous eras, diseases spread more slowly. As transportation became faster and more frequent, diseases began to spread more quickly. The Black Plague, for example, was spread to Europe by merchant ships. In a "globalized" world with fast and frequent travel (especially air travel), pandemics become a problem, as without travel restrictions a virus could reach every continent relatively quickly, a fact pointed out years ahead of the COVID-19 pandemic by Nassim Taleb (*Black Swan* 317). Thus, we can again see how the interdependence of the components of human civilization, and the frequency of interactions between them, leads to a situation in which a single event can have massive consequences (fat-tailed).

Climate Change: The crude technology of earlier times put limits on the extent and level of pollution. Because environmental damage was localized, local actors could be incentivized to reduce or eliminate it. Additionally, the small extent of environmental damage meant that different locales were relatively environmentally independent—the smog of Manchester had little to no effect on the skies of Texas.

Advances in technology have increased the scope of environmental damage. This is especially clear in the case of anthropogenic climate change. CO2 emissions have an impact on the entire planet, no matter where they originate. The planet is now a system of interdependent components. Additionally, since environmental damage is not localized, local actors have fewer incentives to reduce or eliminate it.

These are merely a few examples of global civilization's increases in complexity and the attendant risks. More could be given—for example, the increase in the extent and speed of immigration (and the cultural changes it brings), or the many changes brought about by the internet.

4.2 Best Practices for Governments

We can see from these examples how the increasing complexity of human civilization in the last two centuries has led to novel situations that would not occur in the non-complex environment of earlier human history. These situations could not have been anticipated by the "great names" of political philosophy, such as Plato, Hobbes, Locke, or Kant, because human civilization was not a (highly) complex system during their lifetime. Because these situations are novel, and not directly analogous to situations experienced in the past, our political philosophy needs to be informed by our understanding of contemporary human civilization as a complex system. This does not imply that one can derive a complete political philosophy from complex systems, because political philosophy involves normative considerations that cannot be derived from scientific findings. Complex systems science cannot, for example, tell us anything about the existence and nature of rights. On the other hand complex systems science can influence what direction our political philosophy takes, in the way that empirical psychology influences epistemology and neurology influences philosophy of mind. Much in the way that the findings of neuroscience suggest (for instance) the falsity of a crude kind of mind-body dualism, the findings of complex systems science can provide some support to some political philosophies and put the burden of proof on others, although they can never be used to conclusively "prove" any given political philosophy. And similarly, previous work done in political philosophy can help provide a framework within which to integrate the findings of complexity science.

As we have seen, the massive increase in the complexity of human civilization continues to have both positive and negative effects. It is a truism to say that we want our governments to be in a place to reap the benefits and avoid the consequences of this negative complexity. More specifically, there are at least two goals related to our already existing governments operating in a complex environment that most people of goodwill can agree on. These goals will be amenable to both "liberals" and "conservatives" alike.

First, we want our governments to be adaptable. A system is adaptable if it can respond effectively to the changes in its environment. As Bar-Yam writes, "If the complexity of [environmental] demands exceed the complexity of a system, the system will fail. Thus, those systems that survive must have a complexity sufficiently large to respond to the complexity of environmental demands" (*Dynamics* 808). Since contemporary civilization is complex, we will need to ensure that our governments can be sufficiently complex as well. As much as possible, we will want to reduce the complexity of the environment by operating at a small scale, so to speak. Due to the tradeoff between complexity and scale, a government can be more adaptable at a smaller scale. At a small scale, we will want to foster "non-toxic variety" in the form of policy innovation. A policy "is a law, regulation, procedure, administrative action, incentive, or voluntary practice of governments and other institutions" ("Definition of Policy"); as new situations emerge from the complex environment, new policies will be needed to deal with these situations. Although this goal seems more amenable to the "liberal," conservatives will have no problem with policy innovation in general, provided that it proceeds in a gradual manner, builds upon previous advances, and respects the autonomy and freedom of citizens.

It is not enough to merely increase the range of potential responses to the environment—our governments also need means of filtering out good and bad responses to the environment. A maladaptive response to the environment will be just as bad as no response at all. Thus, our second goal is to filter out *toxic policies*. A policy can be toxic in four ways. First, it can unintentionally cause grave moral harm[3] or easy opportunities for such harm. An example would be "stop-and-frisk" policing implemented by New York City in the early 2000s. Research by the ACLU of New York revealed that the majority of those stopped by the police were Black or Hispanic; moreover, the majority of those stopped were innocent ("Stop and Frisk Data"). Similarly, in 2011 force was used against Black and Hispanic people who were stopped and frisked far more often than it was used against White people ("2011 NYPD Stop and Frisk Statistics). Multiple studies, including a Washington Post analysis, revealed that the program had a negligible effect on rates of crime (Keating and Stevens). Whether or not the policy was intended to negatively impact minorities, it opened the door to racial profiling and brutality against minorities.

Second, a policy can be toxic by failing to achieve what it set outs to achieve. For example, the D.A.R.E. program was instituted to prevent "alcohol, tobacco, and illicit drug use among school-aged youths" (West and O'Neal). However, a meta-analysis of studies on D.A.R.E. by West and O'Neal showed that "Project D.A.R.E. is ineffective…the effects [of the program] did not differ significantly from the variation one would expect by chance." Because

[3] It is beyond the scope of this thesis to precisely delineate the boundaries of "grave moral harm." Nevertheless, it is easy to distinguish between obvious instances of grave moral harm (e.g. a policeman raping and murdering an innocent civilian) and trivial moral harm (a DMV employee being rude to me as I am renewing my driver's license), and it is these obvious instances of grave moral harm that should be kept in mind when this phrase is used in the thesis.

D.A.R.E. failed to achieve its goals, yet took up valuable time and resources, it is a toxic policy. Similarly, "stop and frisk" is a toxic policy because, in addition to creating opportunities for grave moral harm, it failed to significantly reduce crime.

Third, a policy can be toxic by achieving the opposite of what it sets out to achieve. In another study of the D.A.R.E. program, Rosenbaum and Hanson showed that "exposure to supplemental drug education [other than D.A.R.E.] appears to have been largely counterproductive" and resulted in "significantly greater negative attitudes toward police, more positive attitudes toward drugs, alcohol, and cigarettes, and more delinquency" (399) The supplemental drug education programs studied by Rosenbaum and Hanson would thus be an example of a toxic policy.

Fourth, a policy can be toxic by generating other negative unintended consequences (regardless of whether the goals of the policy obtain). For example, Nicholas Kristof noted that some parents of special-needs children pulled those children out of school in order to continue receiving checks from the Supplemental Security Income program—if the children stayed in school and improved academically, the family could lose their eligibility for the checks ("Profiting From a Child's Illiteracy). This is an example of a toxic policy—even if it achieved its goal of providing financial assistance to families with special-needs children, it created the unintended consequence of promoting illiteracy and lack of education. This lack of literacy can in turn cause further and greater problems as the children grow up—the "cascading effects" discussed earlier in the thesis. If the unintended cascading effects of a policy reach the systemic level, they could have highly damaging or even catastrophic consequences.

Thus, governments need filtering mechanisms in place to make sure that toxic policies are not put in place—and if they are put in place, that they will be initially implemented at a small-scale, and quickly vetoed or modified. Although the goal of stifling toxic policies seems more amenable to the "conservative," it is in the best interest of liberals to attain this goal as well; if one thinks the government should take on a wide range of responsibilities, then one will naturally want the government to fulfill all those responsibilities in the best possible manner.

The overall meta-goal of these two goals is to allow governments approximate the evolutionary environment of naturally-occurring complex systems. To do that, we will apply the best practices of complex systems engineering listed above: 1) Foster non-toxic variety at a small scale, 2) expose to the ecological early, 3) detect and fail the toxic fast and locally, 4) allow for means of communication between parts of the system to copy successful changes, and 5) resist the temptation to quickly scale solutions. However, an important caveat is in order. Although we want to approximate the evolutionary environment in some respects, political enterprises are goal-directed and take into account normative considerations, unlike the unguided, amoral process of biological. The approach recommended in this thesis is most definitely not a form of "Social Darwinism"--I take it for granted that governments should not inflict grave moral harm upon their citizens, whether intentionally or unintentionally. Thus, our best practices are not inviolable rules; they will sometimes need to be balanced against other considerations, including considerations of justice or rights. Some of these considerations will be addressed in section 5.1. Different political philosophies may make different tradeoffs. This thesis is not meant to be a presentation of a complete political philosophy, but merely an illustration of a smaller part of political philosophy. And although, as it will be argued, complexity science points in the direction of Anglo-American conservatism, the modest conclusions of this thesis could easily be fit within a more left-leaning framework.

In the next sections of this thesis, I will argue that the goals of government in a complex environment can best be met by adhering to the following maxim: The burden of proof is on the one who proposes that jurisdiction should be larger rather than smaller, power more centralized rather than decentralized, policy making more top-down than bottom-up and policy implementation more uniform rather than discretionary, and jurisdiction larger rather than smaller—or in other words, on the one who proposes an increase in the scope of government. I will make the case that the different claims in this maxim are applications of the aforementioned best practices of complex systems engineering, and that they help us achieve the goals of government operating in a complex environment. After I have done that, I will discuss situations in which the burden of proof is passed—situations in which there is good reason for increasing the scope of government.

One quick note is worth making: The remaining discussion makes use of many examples. These examples are merely illustrative—they are not attempts to "prove" the claims being made. This is for two reasons. First, the claims being made are derived from the nature of complex systems, not from inductive study. If these claims were derived from inductive study, the small-scale, anecdotal nature of the examples would not be even close to sufficient to prove the claims. Second, because the best practices are not absolute rules, but are always balanced against other considerations, it is always possible to cherry-pick examples which seem to confirm or disconfirm them. However, neither kind of example is sufficient by itself to vindicate or discredit the best practices.

4.3 Small-scale Vs Large-scale Jurisdictions

Claim: The burden of proof is on the one proposing an increase in the size of a jurisdiction.

Treisman defines "jurisdiction" in the following way:

> Each state governs a particular territory, or set of points in space. A *jurisdiction* is a subset of the territory consisting of contiguous points. While some jurisdictions contain others, I assume the borders of jurisdictions do not cross the largest jurisdiction comprises the state's entire territory, and I call this the *first tier jurisdiction.* (22)

Recall that the emergent properties of complex systems are scale-variant—they change as the size of the system changes. The more interdependent parts, the more complex the system is. Thus, larger jurisdictions will be more complex, and thus have qualitatively different properties than smaller jurisdictions. Limiting the size of jurisdictions limits the complexity of the environment. There are four benefits that can come from limiting the size of jurisdictions.

First, toxic policies enacted in smaller jurisdictions have smaller negative impacts than those enacted in larger jurisdictions. This is true for two reasons. First, fewer people are affected by toxic policies in a smaller jurisdiction. Second, because increases in the size of a complex system can lead to unpredictable emergent properties, some policies may have unintended consequences whose effects become non-linearly greater (the "cascading" effects discussed earlier) as the scale of the policy increases. These effects could cause systemic, rather than local harm, within a larger jurisdiction.

Second, smaller jurisdictions have shorter feedback loop. Assuming feedback mechanisms (such as voting and petitions) are in place, residents of a smaller jurisdiction will be more easily able to transmit feedback to its decision-making bodies, due to several factors: the smaller quantity of feedback, lower communication costs, and fewer layers for the feedback to travel through. Thus, smaller jurisdictions are more easily able to quickly copy successful changes and fail unsuccessful ones.

Third, co-ordination of complex activities is easier in smaller systems. This is due to the tradeoff between complexity and scale mentioned earlier. For example, in a conventional military unit, one officer controls a group of soldiers. It is relatively easy to coordinate a large number of soldiers in performing simple tasks, such as marching in formation. However, as the behavior of the units grows more complex, an instability develops because “a group of individuals whose collective behavior is controlled by a single individual cannot behave in a more complex way than the individual who is exercising the control” (*Dynamics* 808). Beyond a certain point, increasing the complexity of the group’s actions requires giving more autonomy to members of the group. Trying to coordinate a large number of people in performing a highly complex set of actions is beyond the scope of human abilities; however, it is possible to coordinate more complex actions at a smaller scale (e.g. fewer people).

Another example of this comes from economics. Markets do a much better job of resource allocation than centralized control. As Bar-Yam points out,

> For an individual to allocate the supply of oil, all of the needs of different users in amounts and times, the capabilities of different suppliers, and the transportation and storage available must be taken into account. Even if one were to suggest that a computer program might perform the allocation, which is recognized as a formally difficult computational problem, the input and output of data would often eliminate this possibility (*Dynamics* 813-814).

Markets are able to perform what centralized control cannot—complex resource allocation on a large scale. However, it is important to note that on a small scale, resource allocation is possible. Although a centralized controller can’t optimally allocate resources for an entire country, individuals and organizations do relatively well at allocating resources for themselves. Similarly, although the small teams used by modern militaries grant more autonomy to individual structures, they still have a well-defined command structure that allows them to coordinate in pursuit of a common goal. Coordination becomes more difficult as the amount of actors or the complexity of the tasks to be performed increases. One must reduce the amount of actors being coordinated, grant the actors more autonomy, or both.

Reducing the size of jurisdictions—trading scale for complexity—allows for more effective coordination. This is especially useful in initiatives that require large compliance on the part of the population. For example, as of February 10, 2021, Bhutan had only suffered one death from the novel Coronavirus (Drexler). The measures taken to ensure such results included closing public venues, quarantining possibly infected people, contact tracing, several lockdowns, virus testing and providing resources to those whose livelihoods were affected by the virus. These measures required both coordination on part of the service deliverers and cooperation on the part of the populace, and this was relatively easy to achieve, in part, because of Bhutan's small size—it has a population of around 750,000 residents (National Statistics Bureau) and the area of the country is roughly 18,000 square miles ("Geography"). Coordinating collective action on this scale is easier than coordinating it on the scale of a country the size of the United States. Note that our argument here is merely that coordinating complex actions on a small scale is *in general* easier. There may be other obstacles to successful coordination other than scale (e.g. extreme polarization).

Finally, the tradeoff between complexity and scale means that smaller systems are able to implement a wider variety of environmental responses than larger systems. As Bar-Yam writes, "for any system, there is a fundamental tradeoff between the number of behaviors a system can have and the scale of those behaviors" ("Introduction" 4). Since a smaller system can manifest a larger amount of behaviors, it will necessarily be more adaptable. All things being equal, smaller jurisdictions will be more adaptable than larger ones.

Many nations today have a nested structure made up of multiple levels of jurisdiction. For example, the United States has within it lower tiers of jurisdiction, such as states, counties, and cities. In such a nested structure, it is possible to "split the difference" in regards to situations where it is necessary to take action at a larger scale (situations that will be addressed in section 5.1). A government can delegate many responsibilities to local governments (small jurisdiction), but reserve some for the national government (large jurisdiction). This brings up the question of how power shall be arranged between the nested jurisdictions. Unsurprisingly, complex systems science can suggest an answer to this question.

4.4 Centralization vs. Decentralization

Claim: The burden of proof is on the one proposing a move from a decentralized manner of organization to a centralized one.

Centralization refers to the "distribution of power between a centre and the outlying parts" (Bealey 45). Treisman describes a hypothetical completely centralized government in the following way:

> A single government exists, based in the nation's capital, with the whole national territory as its jurisdiction. This government directly chooses all public policies for all parts of the territory and implements and adjudicates them itself. It can, of course, implement different policies in different parts of the territory if it so chooses. (22-23).

Decentralization occurs when power is distributed from the central government to various subunits. Scholars of decentralization recognize (at least) three kinds of decentralization, administrative, political, and fiscal. The arguments made in this thesis will primarily focus on political decentralization. Political decentralization "is a set of institutions (e.g. constitutional and electoral reforms) designed to devolve political authority, especially electoral capacities, to subnational actors" (Grossman 49). Grossman gives the example of "the popular election of governors, mayors, or village heads who previously were appointed to their positions" (40). In a politically decentralized system, lower tier jurisdictions have "some decision making authority […] that is difficult to reverse" or have "some rights to select lower-level officials, or both" (Treismann 23). In short, power is distributed throughout a decentralized system, rather than concentrated in the central government.

There are four reasons for preferring political decentralization. First, politically decentralized subunits can craft their own unique policies, allowing lower-tier jurisdictions to act as "policy laboratories" and thus approximate the evolutionary environment. Second, politically decentralized subunits can quickly reject toxic policies, thus keeping them from spreading. Third, a collection of semi-autonomous subunits can organically "evolve" policies and scale them slowly. Fourth, small-scale coordination of the kind described above is more easily accomplished in units that are politically decentralized. Let us address each of these reasons in detail.

First, politically decentralized subunits can craft their own unique policies. The usual, "axiomatic" reason for preferring this is that "local laws can be adapted to local conditions and local tastes, while a national government must take a uniform—and hence less desirable—approach. [...because] more people can be satisfied by decentralized decision making than by a single national authority" (McConnell 1493). For example, a location with frequent hurricanes or tornadoes has a reason to enact certain laws governing the construction of buildings that would be superfluous in a different environment. A densely populated city like New York City will benefit from different laws than a sparsely populated county in rural Montana.

However, another benefit that is related to complex systems theory is that decentralized subunits can act as "policy laboratories" testing policies as "pilot programs"—and thus foster variety and policy innovation at a small scale. If each subunit in a nation can craft its own laws (within constitutional limits) the range of possible responses on the part of the government is greatly increased. A system with more opportunities for policy making will be more adaptable than one with fewer opportunities.

The link between political decentralization and policy experimentation should not be overstated. For example, as Treisman points out:

> [C]entral governments in centralized states can enact different policies in different localities, and do so all the time. Moreover, the governments of politically centralized states certainly can and do conduct localized policy experiments. This is true of centralized dictatorships. In the Soviet Union, Leonid Brezhnev authorized regional economic experiments, extending successful ones to other areas, and so did both Mao and his successors in China. It is also true of centralized democracies. In the United Kingdom—usually considered among the most centralized—the government routinely tests policies locally before 'rolling them out' nationwide. One 2003 survey identified 'well over 100' such pilot schemes conducted in the previous five years and even worried that central authorities might run out of test sites" (Treisman 224).

Thus, large-scale political decentralization might not be necessary to achieve "non-toxic variety"—administrative decentralization coupled with incentives for experimentation would be able to achieve the same goal.

However, political decentralization is necessary for the rejection of toxic policies. If subunits have some level of veto power, they can keep toxic policies from spreading and creating negative effects on a grand scale. While higher levels of government can certainly monitor the effects of policies and alter or abolish them as needed, subunits are uniquely qualified to do so for two reasons. First, those who live in the subunits have more firsthand knowledge of the effects of policies than those who do not. If a policy crafted in Washington D.C. has negative effects in Fargo, North Dakota, the government officials in Fargo will know about it before the policy crafters in D.C. do. Second, those who live in the subunits are often the first to feel the effects of a given policy. The residents of the subunits have "skin in the game," while those who craft the policy do not. In both cases, the residents of the subunits have relevant experience of the policy that allows them to identify it as toxic. Thus, subunits must have the power to veto policies so that they can quickly get rid of toxic policies, not only in order to prevent an unjust imbalance of power, but also to fail these policies quickly and locally.

A third reason for preferring decentralization is that it can allow successful policies to "evolve," scaling them upward slowly in a trial-and-error fashion. Recall that markets can perform the complex task of resource allocation without a central controller. Similarly, a collection of decentralized units can perform the task of policy innovation without a central controller. Allowing each subunit freedom to create, implement, and reject policies can allow the subunits to purse a trial-and-error approach, copying the successful policies of other subunits and rejecting or modifying policies as needed. In this way, a collection of decentralized subunits can scale solutions upward slowly, with modifications at every step of the process instead of being rolled out to the entire nation at once. For example, same-sex marriage was initially legalized in Massachusetts in 2003 (Burge). Other states gradually adopted similar legislation, so that by the time that same-sex marriage was officially established by the Supreme Court in *Obergfell vs. Hodges* in 2015, a majority of states had already passed legislation establishing same-sex marriage (Denniston).[4] Similarly, although marijuana is still a Schedule I controlled substance under federal law ("Controlled Substance Schedules"), many states, beginning with Washington and Colorado, have experimented with marijuana legalization (Coffman and Neroulias). As of now, Idaho, Nebraska, and Kansas are the only states without any "public cannabis access program," while eleven other states have a "CBD/Low THC" program (Harmann). If these

[4] To some degree this example is ambiguous—as I shall argue later, civil rights are best handled at the highest level of government, then this example may not work if one considers marriage one of those civil rights. However, the example does work if one does not consider marriage a civil right, or is agnostic about the status of marriage.

programs continue to be successful at a small scale, they can be scaled upward slowly and modified over time. In this way, a collective of decentralized subunits can approximate the evolutionary environment—policies can "evolve" through trial-and-error, and the policies that pass this "test" can eventually be adopted by all the subunits, while the policies that fail the test can be discarded.

A fourth reason for preferring decentralization is that, as has been previously noted, coordination of complex activities is easier on a small scale than a large scale, especially when these projects require large-scale compliance on the part of the population. Giving some autonomy to lower tiers (e.g. state and local governments) allows them to coordinate these activities as opposed to coordinating them on a national scale. While it is possible for the federal government, for example, to coordinate many small-scale actions, it is likely easier for state and local governments to do so.

Although we have seen the benefits of smaller jurisdiction size and decentralization, there are many cases in which the best course of action is intervention a central authority at the largest jurisdiction size. These will be discussed in detail in section 5.1.

4.5 Top-down vs. Bottom-up Policy Making

Claim: Policy making should be conducted in a bottom-up rather than top-down fashion as much as possible. The burden of proof is on the one proposing a move from a bottom-up to a top-down process.

By policy making I mean any attempts to create, alter, abolish, or institutionalize a policy. This loose definition includes actions such as citizens voting in a referendum, lawmakers negotiating the passage of a bill, and judges declaring a law to be constitutional. Though there are important differences between these different types of actions, they all involve a (proposed) change to policy, even if that change is as simple as the official recognition that a policy is constitutional.

An example of maximally top-down policy making would be the writing and implementing of an executive order. An example of maximally bottom-up policy making would be the referendums that are a feature of direct democracy. An example of a mixed policy making would be legislation in Congress. Elected representatives are chosen directly by the citizens of the United States (bottom-up). These representatives in turn craft and modify policies through a process of deliberative democracy that roughly mirrors the same process that put them in office. Finally, the legislation is sent to the desk of the president, who has veto power over it (top-down).

The more "bottom-up" a policy-making process is, the more input those at lower levels have. For example, citizens have the maximal level of input in electing representatives, voting for referendums, and petitioning. They have some level of input in congressional legislation—although they cannot themselves vote for or against congressional legislation, they can elect members of congress and pressure them to vote a certain way. And citizens have no official input on executive orders from the president, aside from electing the president who enacts them.

All things being equal,[5] there three reasons for preferring bottom-up policy making. First, bottom-up policy making approximates the "trial-and-error" evolutionary environment, aggregating preferences and filtering out potentially toxic ideas in a market-like fashion. Second, bottom-up policy making allows toxic policies to be failed locally by those who have skin in the game. Third, the slow nature of bottom-up policy making allows successful policies to be copied

[5] I assume, for example, that the parties involved in the bottom-up processes are reasonably well-informed, that the processes are not marred by corruption, that there is some form of fair access to participation in these processes (e.g. no racial discrimination, etc.). In the real world, structural barriers to well-functioning bottom-up processes exist, and proposing solutions to these barriers is beyond the scope of this thesis.

and scaled in an organic fashion, while preventing toxic policies from being rolled out to the entire nation. I will address each of these reasons in detail.

First, bottom-up policy making approximates the "trial-and-error" evolutionary environment in a positive way. There are two ways in which it does this. First, it allows for greater variety in policy creation. Top-down policy making is typically performed by an individual or relatively small group of individuals—for example, the President, or the Supreme Court. On the other hand, bottom-up policy making can be performed by larger groups of people, including the entire nation. The larger the group of people who are involved in creating potential policies is, the greater the variety of potential policies, and thus the higher the probability of successful policies being created. Or, to put it in the terminology of this thesis, non-toxic variety is fostered, and the range of possible responses by the government is increased.

Second, as well as creating a mechanism for policy generation, bottom-up policy making allows for a filtering mechanism for toxic policies. The majority needed to enact policy change ensures that some (though not all) toxic policies will be rejected. For example, policies crafted to benefit one small interest group will have a more difficult time being instituted if they require the support of a majority to be instituted. The claim here is not that bottom-up policy making reliably filters out all toxic policies—it is unlikely that any filtering mechanism could do so. Rather, it is that bottom-up policy making is necessary to filter out *some* toxic policies which would otherwise be enacted.

A second reason for preferring bottom-up policy making is that it can allow policies to be failed quickly and locally by allowing the input of those who are affected by those policies—the people with "skin in the game." All things being equal, if local communities are affected by toxic policies, the community members have an incentive to quickly alter or abolish them. On the other hand, policy makers in Washington who do not feel the effects of these have no incentive to change or eliminate toxic policies. When local communities, not policy makers in Washington, bear the costs of toxic policies, the members of these communities should have veto power over those policies, unless such veto power would violate constitutional rights or lead to grave moral harm. For example, a community should not be able to veto desegregation laws, for instance. More controversially, they should not be able to veto anti-pollution legislation if that legislation protects other communities from externalities. On the other hand, communities should be able to veto a wide variety of policies that do not involve protection from externalities. In this way, communities will be able to quickly stop the spread of some toxic policies and keep them from causing both systemic and local harm.

A third reason for preferring bottom-up policy making is that the slow nature of bottom-up processes prevents policies from being scaled too quickly. Bottom-up processes, such as deliberative democracy, deliver slower results than top-down processes, primarily because achieving the consensus necessary to pass a law or enact a petition requires a lengthy process of coalition-building. This process must be repeated multiple times if, for example, each subunit decides on a particular policy for itself rather than that policy being mandated to all subunits. This gives time for potential “cracks” in the policy to appear. If there are toxic aspects to the policy, they are more likely to become visible if the policy is both debated and tested over a longer period of time.

A fourth reason for preferring bottom-up policy making is that successful policies can be scaled organically. Bottom-up processes are a means by which successful policy changes can be copied. If a policy proves to be successful in a given subunit, voters in another subunit have a reason to vote for it in their subunit. If a successful policy is tested out in multiple subunits, and continues to be successful (both in terms of non-toxicity and in terms of enjoying voter support), bottom-up processes give us a method it can be scaled organically, passing the "test of time" in a way that minimizes potential side effects. Here we can see a third role of top-down policy making processes: in addition to preventing grave moral harm and solving collective action problems unable to be solved on a lower level, they can also institutionalize policies that have been shown to be successful in an organic fashion, thus preventing further (unhelpful) meddling (perhaps from special interests). Top-down processes (e.g. the ruling of the Supreme Court, an executive order) can act as a kind of "seal of approval" on policies that have already been shown to be non-toxic in multiple contexts. And we can see this process in the previous examples of same-sex marriage and marijuana legalization; in both cases, the policies were scaled organically through bottom-up processes, and in the case of same-sex marriage, the ruling in *Obergfell vs. Hodges* served as a "seal of approval" on a policy that had been shown to have widespread support and no obvious toxic effects.

4.6 Uniformity vs. Discretion

Claim: The burden of proof is on the one proposing a move from a discretionary policy implementation to more uniform policy implementation.

Even if mechanisms for bottom-up policy making are in place, there are still times in which policies require changes that bottom-up mechanisms cannot deliver. The feedback mechanisms of bottom-up policies are slow, and as I have argued, this is one of their advantages. However, this can also be a disadvantage when there are drawbacks to the long amount of time it can take for bottom-up processes to make changes. Additionally, sometimes those with the most accurate information about the nature and effects of policies are not those affected by the policies, but those implementing the policies. It is dubious that the inmate of a maximum-security prison or a psychiatric ward is in a better place to evaluate policy implementation than a guard or a psychiatrist. Here I assume good faith on the part of the policy implementers, as no political system can entirely eliminate the presence of sadistic or corrupt service providers[6]. Good-faith policy implementers need to be given some discretion in implementing policy, for the following two reasons.

First, discretion increases the amount of possible responses to (unforeseen) situations. Recall that systems in a complex environment must have a set of responses that is equal to the complexity of the environment. No policy can specify every possible response to 1) local environments that differ due to material and cultural conditions and 2) an environment that is constantly changing due to a large number of frequent interactions between interdependent components. Allowing policy implementers greater amounts of discretion rather than enforcing a uniform implementation of a policy gives policy implementers a wider range of responses, allowing them to adapt to the environment quickly rather than waiting on bottom-up processes to institute adaptations.

Second, allowing policy implementers discretion may allow policy implementers to prevent a policy from becoming toxic. Discretion can give policy implementers the freedom to avoid actions that would result in toxic effects unintended by the policy makers.

A good example of the benefits of discretion comes from the controversy over mandatory sentences and sentencing guidelines in criminal law, both of which removed discretion from judges in sentencing offenders. While it is obvious how mandatory minimum sentences remove discretion from judges, sentencing guidelines achieved the same effect by creating what Gertner called the "Civil Code Ideology of the Sentencing Reform," where "judges were to be clerks, not

[6] Of course, some systems will be better at monitoring service providers, punishing bad behavior, and incentivizing good behavior.

interpreting the document [the United States Sentencing Guidelines], which after all, was essentially perfect, but simply providing sentencing 'answers.'" (533-534). Decisions made by judges became "formulaic" (536)—judging was reduced to woodenly applying the guidelines to each crime.

Reducing sentencing to the application of a relatively simple decision procedure led to cases in which punishments did not fit the crime. For example, Forer relates the case of an offender named Michael, who held up a taxi driver using a toy gun after he had lost his job. In Michael's case, there were extenuating circumstances that the judge—in this case, Forer—needed to take into consideration: "[T]his was a first offense, no one was harmed, Michael acted under the pressures of unemployment and need, and he seemed truly contrite […] he posed no danger to the public." Forer ruled the minimum sentence unconstitutional and sentenced "Michael to 11-and-a-half months in the county jail" and "two years' probation following his imprisonment conditioned upon repayment of the [stolen] $50." However, this decision was appealed and Michael was "resentence[d] […] to a minimum of five years in the state penitentiary."

Apart from the possibility of grave moral harm involved in meting out a disproportionate sentence to an offender, the policy of mandatory sentencing was toxic because it undermined the goals of the policy, and of the justice system as a whole. As Forer writes,

> The usual grounds for imprisonment are retribution, deterrence, and rehabilitation. Michael had paid his retribution by a short term of imprisonment and by making restitution to the victims. He had been effectively deterred from committing future crimes. And by any measurable standard he had been rehabilitated. There was no social or criminological justification for sending him back to prison.

Similarly, if sentencing "guidelines" become de facto mandatory rules for sentencing, they can prevent judges from exercising the discretion necessary to achieve the goals of "retribution, deterrence, and rehabilitation."

On the other hand, it is clear that there should be some limits on discretion. Beyond the obvious cases such as bribery, corruption, and sadism, it seems that there are at least two limits on discretion. First, discretion should not be allowed if well-intended modifications to the policy result or would result in failing to achieve policy goals or achieving of the opposite of what the policy sets out to achieve. Second, discretion should not be allowed if modifications to the policy would significantly change the goals or outcomes of the policy. In both cases, I assume that the initial policy is non-toxic.

For example, Supreme Court Justice Anthony Kennedy pointed out in a speech to the ABA that "Before [sentencing guidelines] were in place, a wide disparity existed among the sentences given by different judges, and even among sentences given by a single judge." Without *some* consistency on the part of judges, sentencing will not consistently achieve the goals of "retribution, deterrence, and rehabilitation," and may even create unintended consequences, such as increasing the rate of crime—thus, achieving the opposite of what the policies involved set out to achieve. Similarly, if policies regarding some drug-related offenses are intended to lead to rehabilitation, but allowing judges discretion results in an overall approach of deterrence instead, discretion is changing the goals of this policy, and should not be allowed. These kinds of changes to policy should be made formally rather than informally, because they do not involve small-scale modifications of the means, but large-scale changes to the ends of the policy.

With some discretion allowed, policy implementation will necessarily be a trial-and-error process. It is important to ensure that the errors are quickly and decisively rectified, and the successes replicated. Thus, two safeguards may be necessary. First, there should be some sort of system in place for prohibiting modifications that are known to create toxic results. This should include a mechanism for formalizing or institutionalizing new prohibitions on certain modifications, as it will not be possible to know in advance whether some modifications will have toxic results. Second, there should be a similar mechanism in place for formalizing or institutionalizing modifications that have been repeatedly shown to have positive results in multiple contexts. Perhaps some sort of bottom-up process would be appropriate here.

5. Objections and Discussion

5.1 Meeting the Burden of Proof

As stated before, these best practices must be balanced against other considerations. I have argued in this thesis that the burden of proof is on the one proposing an "upward"[7] move. I will now present two situations in which an upward move meets the burden of proof, and thus the best practices may not apply. These situations are not exhaustive—there may be other instances in which considerations other than adaptability force us to abandon or modify one of the best practices.

The first situation in which the burden of proof is met is when allowing a jurisdiction to be small, power decentralized, policy making bottom-up, or policy implementation discretionary, causes or creates easy opportunities for grave moral harm. One example of such moral harm is the violation of rights. For example, if human rights are universal, then they should be enforced "across the board." Allowing different locales autonomy in deciding who is allowed to have human rights would result in human rights violations. The rights to freedom of speech, freedom of religion, and freedom from slavery are enshrined in the United States Constitution (Amend. I

[7] For the purposes of this paper, I will refer to increasing the size of jurisdictions, centralizing power, making a process more top-down, and removing discretion/increasing uniformity as "upward" moves and their opposites as "downward" moves.

and XIV); making these rights the subject of bottom-up policy making and experimentation would almost certainly result in grave moral harm.

Additionally, Calabresi and Bickford argue that "national government is better at handling civil rights issues than are the states" (132), because the larger scope of the national government means that groups that would be oppressive majorities on a smaller scale have less of a share of power and influence. It is easier for the Ku Klux Klan to influence policy in a small city than in the entire nation, where they must compete with other interests. Thus, the burden of proof is passed in the case of the enforcement and enumeration of rights.

Another example of grave moral harm is externalities, where (part of) the cost of an action is borne by those who do not commit the action. The lack of "skin in the game" on the part of the actors means that a third party is necessary to monitor the actors and enforce certain behaviors. Calabresi and Bickford use the example of a factory in one state polluting another state. They write: "If a state could realize the economic benefits of a factory while the costs of its pollution fell mostly on other states, the polluting state would have no incentive to clean up its act. National regulation of clean air and water is thus essential to correct for the externalities problem" (132). Expanding the scope of government is necessary in these scenarios where "skin in the game" does not provide an incentive to correct action.

An extreme example of an externality is a *ruin problem* (Taleb et al., "The Precautionary Principle"). A ruin problem is a collective action problem with a risk of an infinite or near-infinite loss. In these kinds of problems, "losing" makes it impossible for everyone involved to continue playing the game. Thus, one cannot use repeated experience to determine the correct means to solve this problem, because any errors in the trial-and-error process would end the process. Ruin problems involve externalities because the "losing" behavior can be performed by only a few agents, but the entire system is destroyed by that behavior. In these situations, the best practice is a comprehensive prohibition of the risk-bearing behavior, coming from the highest levels (centralized, top-down, uniform, and at the level of the largest jurisdiction).

The most pressing ruin problems today are anthropogenic climate change, pandemics, and nuclear war (Taleb, et al., “Precautionary Principle” 9-10). Each of these is a ruin problem because in each case irreversible systemic damage on the level of human civilization is a possible outcome. For example, possible outcomes of anthropogenic climate change include ecocide or mass extinctions—both of which cannot be recovered from in the lifetime of any human being. Even if massive deaths from pandemics, climate change, or nuclear war have not occurred in the past, the possibility of systemic, irreparable damage from these events should change the way we prepare for them. Dealing with them in the trial-and-error, decentralized manner that this thesis generally advocates is like playing Russian roulette with other people’s lives. Even if the chances of these negative outcomes are small or uncertain, the possibility of total system failure changes the strategy that needs to be taken. Instead of a trial-and-error approach that this thesis recommends to tackling most policy problems, the approach should be “minimize all errors.” Allowing subunits and individual actors to have discretion and input to climate policy is like allowing discretion and input to a person tasked with planting a bomb. Just as it would be irresponsible to allow citizens to decide whether to allow slavery in a given subunit, it is irresponsible to allow citizens to choose policies that we know increase the risks of anthropogenic climate change, and consequently the risks of ruin to human civilization. In both cases, grave moral harm results. Thus, limitations on actions that can exacerbate each problem should be enacted and enforced by the highest level of government in the most uniform way.

This does not, however, imply that solutions to each problem should be enacted in the same way. A large-scale, centrally-enforced, uniform policy limiting risky behavior is necessary, but this does not imply that there should be similar policies mandating non-risky behavior. This is because prohibiting non-risky behaviors limits potential responses to the environment. For example, limiting emissions—and thus, placing limits on the use of energy sources such as coal—is consistent with this approach to ruin problems. However, this does not imply that the government should mandate a specific clean energy alternative. Letting the market, and/or state and local governments choose between, for example, solar, wind, nuclear[8], or some combination of the three, will allow the most efficient kind of energy source to be adapted (or, at the very least, will filter out some of least efficient sources). Similarly, a cap-and-trade approach to emissions allows for companies to experiment with different arrangements of emissions while still reducing emissions overall. As long as the behavior that creates risk—in this case, risk of climate change, and therefore ecocide or mass extinction—is prohibited or curtailed, there should be leeway for experimentation with non-risky behaviors. This kind of approach has been shown to be effective in the past. For example, the problem of the hole in the ozone layer was solved in large part by the unilateral adoption of the Montreal Protocol by every member of the United Nations ("About Montreal Protocol"). The Protocol has been a success: "To date, the Parties to the Protocol have phased out 98% of ODS globally compared to 1990 levels" ("About Montreal Protocol")[9]. Similarly, the Acid Rain Program, put in place by the 1990 Clean Air Act Amendments, resulted in a 94% reduction of sulfur dioxide emissions from 1990 to 2019 ("Acid

---

[8] Assuming nuclear plants are not too close together, nuclear power does not constitute a ruin problem. The meltdown of a nuclear power plant is, of course, disastrous and highly dangerous, but it does not cause the failure of an entire system. Fukushima and Chernobyl did not destroy Japan and Russia.

[9] Interestingly, the Republican President Ronald Reagan advocated for the ratification of the treaty—thus creating a precedent for such actions among those who lean right (see Glass).

Rain Program"). In both cases, markets adapted to provide alternatives to the prohibited behaviors.

The second scenario in which the burden of proof is met is a case in which experience has shown repeatedly that a given problem P cannot be effectively solved by bottom-up policy making, decentralized authorities, etc. Therefore, an "upward move" is necessary to solve this problem.

There are several ways that this can play out. First, as Calabresi and Bickford point out, "sometimes there are substantial economies of scale in undertaking an activity or financing a program only once rather than fifty times" (131) In this kind of scenario, collective action to solve P is possible, but the economic costs of doing so on a small scale make it so that P is not solved *effectively*. Second, the (non-monetary) costs of collective action make P impossible or difficult to solve on a smaller scale. Calabresi and Bickford give the following example:

> It would be very time-consuming and expensive for the fifty states to act collectively on foreign policy, or defense, or national economic policy. Some states might refuse to join in policies that a majority of states representing a majority of the people endorse. Such holdout states might trigger a race to the bottom and cause the legal standard of the most permissive state to force all other states to comply, even if a majority of the nation wished otherwise (132).

Third, there are cases in which actors do not have an incentive to cooperate and solve P—think, for example, of "tragedy of the commons" situations. In this case, a third party, such as the government, or a higher level of government, is necessary to ensure that P is solved. To some degree, these may overlap with the externalities discussed above, leading to situations where the government must intervene both to solve an otherwise insoluble problem, and to avoid grave moral harm.

Even if there seems to be a structural barrier to solving problem P without an "upward move," one must be cautious and open to the possibility of creative solutions from below. Just because it is possible to construct a model where "rational actors" solve or fail to solve P does not imply that this model will accurately reflect reality. For example, in the type of collective action problem known as the "tragedy of the commons," rational actors in the models fail to solve the problem. This has led some scholars to conclude that "tragedy of the commons"-style situations cannot be solved except through the intervention of a third party with coercive power, such as the government (Ostrom). However, as Elinor Ostrom has demonstrated, at least some real-life instances of the "tragedy of the commons" have successfully been solved without government intervention. One example she gives is a fishery in Alanya, Turkey. The fishery was a common resource that every fisher had incentive to exploit, thus leading to a "tragedy of the commons." However, the fishers were able to solve this problem, not through government intervention, but through coming together and instituting a set of rules enforced by the fishers themselves—a sort of Social Contract without a Leviathan. The mere existence of a structural problem did not imply that increasing the scope of the government was the only viable solution.

Similarly, one could easily cherry-pick examples of failures to solve P and use them as "evidence" that P cannot be solved without an upward move, when in fact an untried solution that does not involve an upward move exists. Thus, *repeated* experience is necessary to show that an increase in the scope of government is necessary. Of course, this does not apply in the cases of ruin problems discussed above, because one cannot continue the trial-and-error approach without risking systemic destruction. However, in cases where there is no risk of systemic ruin or grave moral harm, we should be cautious about quickly jumping to a solution "from above."

5.2 Potential Objections and Points of Connection

The previous application of complex systems science to political philosophy may seem to be "true but trivial." After all, one may reason, many of the points made—the instrumental value of democracy, the benefits of decentralization, the superiority of discretion ("bottom-up" policy implementation) to uniformity—have been previously argued for in the literature of political philosophy and political science. If the points made in this thesis are uncontroversial, why make them at all?

There are two responses to be made here. The first to admit that, yes, some of the conclusions of this thesis match up with commonly held intuitions about politics. However, the thesis gives a new and different justification of these intuitions that is rooted in how the world works. An analogy might be made with ethics. Ethical projects of the kind undertaken by, for example, Aristotle, Aquinas, Hobbes, and Kant, did not seek to completely overturn common ethical judgments in favor of completely novel conclusions, but instead sought to provide justifications and explanations for widely accepted ethical judgments, even if these justifications and explanations were revisionist in nature (e.g. Hobbes and Kant). Similarly, this thesis provides an explanation for, e.g., the instrumental value of democracy, in terms of complex systems science rather than contingent political realities.

Second, it is not the case that every conclusion of this thesis is true but trivial. While ethical frameworks of Aristotle, Aquinas, Hobbes, and Kant all justified a wide range of commonly held moral intuitions, they also had differing controversial implications (for example, Kant's blanket prohibition on lying). No ethicist will willingly endorse a system that justifies rape and murder, but the different ethical systems that converge on the proscription of rape and murder will diverge at other points. Similarly, while this application of complex systems science to political philosophy justifies some commonly held intuitions, it also lends support to some more controversial ideas in political philosophy.

For one, it helps explicate some right-leaning intuitions about the desirability of "small" government. Taken at face value, this intuition is confusing--does it refer to the physical size of a government's jurisdiction? The amount of interventions made by the government? The size of a government's budget? It is doubtful that those who want a "small" government would be happy with, say, a government that was highly intrusive, but politically decentralized, or a government with a small budget that did not allow any bottom-up input from citizens or subunits. However, as I have shown, complex systems science gives us reasons for preferring small-scale jurisdiction, decentralized authority, bottom-up policy making, and discretionary policy implementation. Thinking about government in terms of these four axes can refine intuitions about the desirability of "small" government while eschewing the unhelpful size metaphor.

More specifically than this, these insights from complex systems science can lend support to the political philosophy of *Anglo-American conservatism*. A word of caution is necessary: it is not the case that complex systems science can "prove" that Anglo-American conservatism is the best form of political philosophy. It merely gives *some* reasons supporting Anglo-American conservatism, and shifts the burden of proof on some issues away from the conservative and onto the "liberal."

The main point of convergence between our complex systems science-informed political philosophy and Anglo-American conservatism is the *trial-and-error* approach to policy innovations. This approach in conservative thought is sometimes referred to as "historical empiricism." Haivry and Hazony, following the seventeenth-century political philosopher John Selden, describe the historical empiricist approach as follows:

> [O]ur reasoning in political and legal matters should be based upon inherited national tradition. This permits the statesman or jurist to overcome the small stock of observation and experience that individuals are able to accumulate during their own lifetimes […] and to take advantage of the "many ages of former experience and observation," which permit us to "accumulate years to us, as if we had lived even from the beginning of time." In other words, by consulting the accumulated experience of the past, we overcome the inherent weakness of individual judgment, bringing to bear the many lifetimes of observation by our forebears, who wrestled with similar questions under diverse conditions.

However, it is not the case that historical empiricism leads to uncritical acceptance of all existing policies. This is true for two reasons. First, it is obvious that inherited traditions, and the policies that have followed from them, can be and have been wrong on many points. Second, changes in the environment may necessitate policy change. For example, it has been argued in this thesis that the increased complexity of the environment in the contemporary world should give rise to policy changes. Anglo-American conservatism is open to political change—its preferred approach to change is one of gradual reform through trial-and-error rather than large-scale reform along "rationalist" lines. Edmund Burke, the "father of conservatism," argued that "temperate reform" was preferable to more radical reform, because it "is permanent, and because it has a principle of growth. He writes,

> Whenever we improve, it is right to leave room for a further improvement. It is right to consider, to look about us, to examine the effect of what we have done. Then we can proceed with confidence, because we can proceed with intelligence. […] A great part, therefore, of my idea of reform is meant to operate gradually: some benefits will come at a nearer, some at a more remove period ("House of Commons").

This is relevantly similar to the small scale trial-and-error approach of complex systems engineering that I have argued for in this thesis. In both cases, politics is viewed as an "ecology" rather than a "top-down engineering project" ("Scala Politica" 3). For both the conservative and the complex systems engineer, "Prudence" rather than rational planning, "is the director, the regulator, the standard" of "the virtues political" (Burke, "Appeal"). Complex systems science grounds this preference for small-scale, trial-and-error changes in scientific insight about the way the world works.

The preferability of the trial-and-error approach does not mean that every policy (institution, etc.) should continually the subject of experimentation. In addition to the aforementioned cases where the trial-and-error approach would lead to moral harm, there are other cases in which the benefits of certain governmental arrangements have been shown through repeated experience. An analogy to science may be helpful. All scientific theories and findings are in principle subject to revision; however, some theories or findings are so well-supported by the evidence that they can be assumed as true and used as the framework for further research. For example, the theory of evolution by natural selection is a well-supported theory; within the framework of this theory, research continues in a trial-and-error fashion. It is in principle possible for a defeater of evolution by natural selection to obtain, but this defeater would have to be overwhelming. Similarly, it is probably not necessary to revise the U.S. Constitution at every step of the way using a trial-and-error method, and in the way that evolution provides a framework for further research, the Constitution provides a framework for further development of policies that can be refined in this trial-and-error way.

Nevertheless, just as seemingly well-supported scientific theories can be overturned by a much-stronger new theory, it is occasionally the case that seemingly well-supported policies or institutions need to be altered or abolished. To this end, Article V of the Constitution provides a mechanism for changing the constitution, but requires three-fourths of the states to assent to an amendment before it can be ratified—thus, large-scale changes to the Constitution can only be made if there is a large-scale consensus in favor of the change (similar to the large-scale consensus that would be needed to *conclusively* overturn a well-established scientific theory).[10] And these large-scale changes may be necessary in order for governments to address the novel situations brought about by the complex environment that is contemporary human civilization.

6. Conclusion

As we have seen, contemporary human civilization is a complex system, made up of multiple interacting, interdependent parts. The properties of human civilization cannot be reduced to the properties of its individual components. Unlike civilization in earlier times, contemporary civilization is fat-tailed; low-probability events can have an outsized impact on the system. These "Black Swan" and "Gray Swan" events can result in cascading effects and even systemic ruin. Surviving and thriving in this complex environment requires that governments be more adaptable than they have in the past. I have argued that the best way to ensure such adaptability in our governments is that certain best practices be followed in order to promote policy innovation while also filtering out toxic policies. Specifically, I made the case that the burden of proof is on the one who proposes that jurisdiction should be larger rather than smaller, power more centralized rather than decentralized, policy making more top-down than bottom-up and policy implementation more uniform rather than discretionary. This set of best practices is weighed against other considerations. Specifically, in cases where following the best practices would result in grave moral harm, or where it has been conclusively shown that a structural problem prevents the best practices from being effective, these practices should be jettisoned. This approach to government is relevantly similar to the tradition of Anglo-American

[10] Obviously, it is the superiority of the new theory in synthesizing and explaining the data that overturns the new theory—but the scientific consensus acts as a kind of "seal of approval" on theories that have been successfully tested. The lack of such a consensus behind scientific ideas such as the denial of anthropogenic climate change or Intelligent Design theory cannot by itself "disprove" these ideas—but it suggests that their explanations of the data are not manifestly superior to conventional explanations.

conservatism, but differs from it, and other previous political philosophies, in its explicit grounding in complex systems science.

The application of complex systems theory to political philosophy is an exciting field; this thesis is merely one attempt at sketching out a program for this application. Many unresolved questions remain. For example, while this thesis has focused on the dangers of "upward moves," such as centralization, it has not addressed how to think about "downward moves" in political systems. In a complex system, "downward moves" will bring about their own qualitative changes in emergent properties—and these moves may be beneficial or harmful. Similarly, it may also be the case that some situations benefit from, for example, decentralization, and some do not (in a manner similar to ruin problems). Perhaps the best approach here is the trial-and-error, historical empiricist approach shared by both Anglo-American conservatives and complex systems engineers.

Similarly, there are questions in political philosophy that complex systems science can never help answer. As mentioned before, complex systems science cannot tell us whether rights exist, whether they are positive or negative, or what these rights consist of. It can, however, guide us in the implementation of these rights. For example, although complex systems science cannot give us any clue as to whether healthcare is a right, if healthcare is a right, then complex systems science *can* give us clues regarding successful healthcare policy—an evolutionary, trial-and-error approach to healthcare policy rather than a "one-size-fits-all," across the board policy.

Similarly, although complex systems science supports some elements of right-leaning political philosophy, it cannot ultimately decide which political philosophy is "best." It may be that, in the last analysis, the burden of proof on large jurisdiction size, etc. is be passed quite often. The trial-and-error approach may ultimately vindicate left-leaning policies that vastly increase the scope of government—the ultimate ironic victory for Anglo-American conservatism. Until that time, the question of political philosophy remains open.

Science is constantly evolving and always subject to revision. Even well-established scientific findings are in principle subject to revision or reinterpretation. No scientist labels himself or herself a "Newtonian," or a "Maxwellian" (Stark 35). Similarly, a political philosophy that is informed by science will change and develop, even if the basic scientific and political framework is tested and found to be solid. It is my hope that this thesis will serve, not as the "last word" on the application of complex systems science to political philosophy, but as an invitation to further research and debate.

Works Cited

"2011 NYPD Stop and Frisk Statistics" *Center for Constitutional Rights*. https://ccrjustice.org/files/CCR-Stop-and-Frisk-Fact-Sheet-2011.pdf. Accessed 29 April 2021.

Akoundi, Ken and John Haugh. "Tail Risk Hedging: A Roadmap for Asset Owners." *Deutsche Bank*. May 2010.

"Assassin's Creed ® Valhalla." *Epic Games*. https://www.epicgames.com/store/en-US/p/assassins-creed-valhalla. Accessed 29 April 2021.

Bar-Yam, Yaneer. *Dynamics of Complex Systems.* Addison-Wesley, 1997.

Bar-Yam, Yaneer, and Alice Ware Davidson. "Environmental Complexity: Information For Human Environment Well-Being." *Unifying Themes in Complex Systems*. Springer, 2006, pp.157-168.

Bar-Yam, Yaneer and Siegenfeld, Alexander. "An Introduction to Complex Systems Science And Its Applications." *Complexity*, vol. 2020, 27 July 2020. https://doi.org./10.1155/2020/6105872. Accessed 29 April 2021.

Bar-Yam, Yaneer. "Introduction to the Strategy and Methods of Complex Systems." *Handbook of Research Methods in Complexity Science.* Edited by Eve Mitleton-Kelly, Alexandros Paraskevas, and Christopher Day, Edward Elgar Publishing, 2018, pp. 5-17

Bar-Yam, Yaneer. *Making Things Work.* Knowledge Press, 2004.

Bealey, Frank W. *The Blackwell Dictionary of Political Science*. Blackwell Publishing, Ltd., 1999.

Betz, Brandy. "TSMC gets large Intel chip order, Apple R&D plant, - report." *Seeking Alpha.* 27 July 2020. https://seekingalpha.com/news/3595175-tsmc-gets-large-intel-chip-order-apple-r-and-d-plant-reports. Accessed 29 April 2021.

Burge, Kathleen. "SJC: Gay marriage legal in Mass." *Boston.com.* Globe Newspaper Company. 18 November 2003. http://archive.boston.com/news/local/massachusetts/articles/2003/11/18/sjc_gay_marriage_legal_in_mass/. Accessed 29 April 2021.

Burke, Edmund. "An Appeal From the New to the Old Whigs." *Complete Works of Edmund Burke.* E-book, Delphi Classics, 2016.

Burke, Edmund. "Speech on Presenting to the House of Commons." *Complete Works of Edmund Burke*. E-book, Delphi Classics, 2016.

"Burst Pipes and Power Outages in Battered Texas." *New York Times.* 21 February 2021. https://www.nytimes.com/live/2021/02/17/us/winter-storm-weather-live Accessed 22 April 2021.

Calabresi, Steven G. and Lucy D. Bickford. U.S. Constitutional Law." *Nomos LV: Federalism and Subsidiarity*, ed. James E. Fleming and Jacob T. Levy, New York University Press, 2014, pp. 123-189.

Coffman, Keith and Nicole Neroulias. "Colorado, Washington first states to legalize recreational pot." *Reuters*. 6 November 2012. https://www.reuters.com/article/us-usa-marijuana-legalization-idUSBRE8A602D20121107. Accessed 29 April 2021.

"Constitution of the United States." *Congress.gov,* https://constitution.congress.gov/constitution/. Accessed 29 April 2021.

Davidson, Donald. "Knowing One's Own Mind." *Proceedings and Addresses of the American Philosophical Association,* vol. 60 no. 3, January 1987, pp. 441-458, http://links.jstor.org/sici?sici=0065972X%28198701%2960%3A3%3C441%3AKOOM%3E2.0.CO%3B2-P. Accessed 31 May 2021.

"Definition of Policy" *CDC* https://www.cdc.gov/policy/analysis/process/definition.html TK

Denniston, Lyle. "Opinion analysis: Marriage now open to same-sex couples." *SCOTUSBlog.* 26 June 2015, https://www.scotusblog.com/2015/06/opinion-analysis-marriage-now-open-to-same-sex-couples/ Accessed 29 April 2021.

Drexler, Madeline. "The Unlikeliest Pandemic Success Story." *The Atlantic.* 10 February 2021. https://www.theatlantic.com/international/archive/2021/02/coronavirus-pandemic-bhutan/617976/. Accessed 29 April 2021.

"Controlled Substance Schedules," *Diversion Control Division*, Drug Enforcement Administration, https://www.deadiversion.usdoj.gov/schedules/. Accessed 30 April 2021.

Duquette, Shane, and Marco Walker-Ng. "How Much Can The Average Man Lift?" *Outlift*. 16 December 2020. https://outlift.com/how-much-can-the-average-man-lift/#4-how-much-can-the-average-man-deadlift. Accessed 29 April 2021.

"Egypt's Suez Canal blocked by huge container ship" BBC. 24 March 2021. https://www.bbc.com/news/world-middle-east-56505413. Accessed 29 April 2021.

Forer, Lois G. "Justice by the numbers; mandatory sentencing drove me from the bench." *Washington Monthly*. 1 April 1992. *The Free Library*. https://www.thefreelibrary.com/Justice+by+the+numbers%3b+mandatory+sentencing+drove+me+from+the+bench.-a012129669. Accessed 26 April 2021.

Gernter, Nancy. “From Omnipotence to Impotence” *Ohio State Journal of Criminal Law*, vol. 4 no. 2, 2007, pp.523-539, https://kb.osu.edu/bitstream/handle/1811/73045/OSJCL_V4N2_523.pdf;sequence=1. Accessed 30 April 2021
.

“Geography.” *Bhutan.com,* 2010, http://www.bhutan.com/country/geography. Accessed 30 April 2021.

Glass, Andrew. “Reagan seeks to protect fragile ozone layer, Dec. 21, 1987.” *Politico*, 21 December 2018, https://www.politico.com/story/2018/12/21/reagan-seeks-to-protect-ozone-layer-1987-1067252. Accessed 30 April 2021.

Grossman, Guy. “Leadership Selection Rules and Decentralized Governance.” *Decentralized Governance and Accountability: Academic Research and the Future of Donor Programming,* edited by Jonathan Rodden and Erik Wibbels, USAID, pp. 49-74.

Grundler, Christopher “The Acid Rain Program: A Success Story” *The EPA Blog*, EPA, 21 May 2020, https://blog.epa.gov/2020/05/21/the-acid-rain-program-a-success-story/. Accessed 30 April 2021.

Haivry, Ofir and Yoram Hazony. “What Is Conservatism?” *American Affairs Journal,* vol. I, no. 2, summer 2017, https://americanaffairsjournal.org/2017/05/what-is-conservatism/. Accessed 30 April 2021

Harmann, Michael. “Cannabis Overview.” *NCSL*, National Conference of State Legislatures, 8 April 2021, https://www.ncsl.org/research/civil-and-criminal-justice/marijuana-overview.aspx. Accessed 30 April 2021.

Hazony, Yoram. “Conservative Democracy.” *First Things*, January 2019, https://www.firstthings.com/article/2019/01/conservative-democracy. Accessed 30 April 2021.

“How do you product test new products?” *McDonalds,* https://www.mcdonalds.com/gb/engb/help/faq/18823-how-do-you-product-test-new-products.html. Accessed 30 April 2019.

Hunt, Ben. “Taiwan Is Now Arrakis.” *Epsilon Theory*, 27 July 2020, https://www.epsilontheory.com/taiwan-is-now-arrakis/. Accessed 30 April 2021.

Jakobsson, Eric, and Working Group 1 Collaborators. “Complex Systems: Why and What?” *New England Complex Systems Institute.* https://necsi.edu/complex-systems-why-and-what?rq=jakobsson. Accessed 20 April 20201.

Keating, Dan and Harry Stevens. "Bloomberg said 'stop and frisk' decreased crime. Data suggests it wasn't a major factor in cutting felonies." *The Washington Post,* 27 February 2020, https://www.washingtonpost.com/nation/2020/02/27/bloomberg-said-stop-frisk-decreased-crime-data-suggests-it-wasnt-major-factor-cutting-felonies/. Accessed 30 April 2021.

Kennedy, Anthony. "Speech at the American Bar Association Annual Meeting." *SupremeCourt.gov,* 14 August 2003, https://www.supremecourt.gov/publicinfo/speeches/viewspeech/sp_08-09-03. Accessed 30 April 2021.

Kristof, Nicholas. "Profiting From a Child's Illiteracy." *The New York Times,* 7 December 2012. https://www.nytimes.com/2012/12/09/opinion/sunday/kristof-profiting-from-a-childs-illiteracy.html. Accessed 9 June 2021.

McConnell, Michael W. Review of "Federalism: Evaluating the Founders' Design" by Raoul Berger. *University of Chicago Law Review*, vol. 54, no., 1987, pp. 1484-1512.

Mitchell, Melanie. *Complexity: A Guided Tour.* Oxford University Press, 2009.

Murray, Charles. *Losing Ground: American Social Policy 1950-1980.* E-book, Basic Books, 2015.

National Commission on Terrorist Attacks Upon the United States. *The 9/11 Commission Report.* Norton Books, 2004.

National Statistics Bureau. https://www.nsb.gov.bt/# Accessed 30 April 2021.

Norman, Joe. "Complex Systems Science: An Informal Overview — Part III: Synthesis of purposeful systems." *Joseph W. Norman,* 28 February 2017. http://jwnorman.com/?p=4134. Accessed 10 May 2020.

Norman, Joe. Complex Systems Science: An Informal Overview — Part I: Reductionism and Emergence." *Joseph W. Norman*, 15 November 2016. http://jwnorman.com/?p=339. Accessed 10 May 2020.

Ostrom, Elinor. *Governing the Commons: The evolution of institutions for collective action.* E-book, Cambridge University Press, 2015.

Remahl, David. "Bellcurve.svg." *Wikimedia Commons*, https://commons.wikimedia.org/wiki/File:Bellcurve.svg. Accessed 31 May 2021.

Rosenbaum, Dennis P. and Gordon S. Hanson. "Assessing the Effects of School-Based Drug Education: A Six-Year, Multi-Level Analysis of Project D.A.R.E." *Journal of Research in Crime and Delinquency*, vol. 35 no. 4, 1998, pp. 381-412, https://doi.org/10.1177/0022427898035004002. 30 April 2021

Stark, Rodney. *Why God?: Explaining Religious Phenomena.* Templeton Press, 2017.

Stevens, Pippa. "The ship that blocked the Suez Canal may be free, but experts warn the supply chain impact could last months" *CNBC,* 29 March 2021, https://www.cnbc.com/2021/03/29/suez-canal-is-moving-but-the-supply-chain-impact-could-last-months.html. Accessed 30 April 2021.

"Stop and Frisk Data." *NYCLU.org* ACLU of New York. https://www.nyclu.org/en/stop-and-frisk-data. Accessed 30 April 2021.

Taleb, et al. "The Precautionary Principle (with Application to the Genetic Modification of Organisms)." Extreme Risk Initiative. *NYU School of Engineering Working Paper Series,* 1-24.

Taleb, Nassim Nicholas. *Antifragile.* Random House, 2012.

Taleb, Nassim Nicholas. *The Black Swan*, second edition, Random House Trade Paperbacks, 2010.

Taleb, Nassim Nicholas. "Scala Politica: Politics and Ethics Under Scaling and Uncertainty." STEM Academic Press.

Taleb, Nassim Nicholas. *Statistical Consequences of Fat Tails.* STEM Academic Press, 2020.

Tate, Ryan. "Google Couldn't Kill 20% Time Even If It Wanted To." *Wired*. 21 August 2013. https://www.wired.com/2013/08/20-percent-time-will-never-die/. Accessed 30 April 2021.

Treisman, Daniel. *The Architecture of Government: Rethinking Political Decentralization.* Cambridge University Press, 2007.

Wescott, Kristin. Out Of The Apollo Masters Fire Emerges The Vinyl Record Manufacturers Association Of North America." *Forbes*, 6 March 2020, https://www.forbes.com/sites/kristinwestcottgrant/2020/03/06/out-of-the-apollo-masters-fire-emerges-the-vinyl-record-manufacturers-association-of-north-america/?sh=1e9fb7cf34bf. Accessed 30 April 2021.

West, Steven L. and Keri K. O'Neal. "Project D.A.R.E. Outcome Effectiveness Revisited." *American Journal of Public Health*, vol. 94, no.6, pp. 1027-1029, https://doi.org/10.2105/AJPH.94.6.1027. Accessed 30 April 2021.

"Undertale." *Steam*. https://store.steampowered.com/app/391540/Undertale/. Accessed 30 April 2021.