Beliefs, Institutions and Development in a Coevolutionary Complex Adaptive System

Bernardo Mueller
Dept. of Economics
University of Brasilia
and
Visiting Scholar at the Ostrom Workshop

October, 2015

> What difference does it make that humans fall short of substantively rational behavior, which would entail full knowledge of all possible contingencies, exhaustive exploration of the decision tree, and a correct mapping between actions, events, and outcomes? The short answer is that it makes a lot of difference. Economic history is a depressing tale of miscalculation leading to famine, starvation and defeat in warfare, death, economic stagnation and decline, and indeed the disappearance of entire civilizations. And even the most casual inspection of today's news suggests that this tale is not purely a historical phenomenon. (North, 2005: 7)

1. Introduction

Twenty-five years ago, in 1990, two books were published within months of each other that would help lay the foundation for a dramatic shift in the way economic research is done. The first book was Douglass North's *Institutions, Institutional Change and Economic Performance* and the second Elinor Ostrom's *Governing the Commons: The Evolution of Institutions for Collective Action*. Together with a few other fundamental works, these two books have helped to bring the role of institutions in most economic research and other social sciences from the background to center stage. In mainstream Economics institutions were typically taken as given and left deep in the background, or were simply ignored. In other social science they were given more consideration, but often the analytic instruments and tools for assessing their impact were lacking. By showing that institutions are crucial determinants of behavior and choices, and that like any other variable, such as prices, production and population, they are susceptible to analysis, these ideas influenced a thriving and vibrant field in the 1990's, and then took the profession by storm in the 2000's. In a remarkable process the concept of institutions moved from the fringes of most fields in Economics and some other social sciences to becoming a fully integrated and transversal concept. In growth theory, for example, Paul Romer, one of the field's central figures, assess as follows where he sees the direction of future research:

> Further out in the horizon, one may hope for a successful conclusion to the ongoing hunt for a simple model of institutional evolution. Combining that with the unified approach to growth outlined here would surely constitute the economics equivalent of a grand unified theory – a worthy goal by which we may be judged when future generations look back fifty years from now … (Jones and Romer, 2010: 242)

Development economics is another notable example. The field has been completely transformed in the past two decades as institutions have been recognized as the fundamental determinant of countries' success or failure to develop. This research agenda, made popular by Acemoglu and Robinson (2006, 2012), has helped to make the concept of institutions respectable to even the most neoclassically inclined researchers and it is now considerably rare to find analyzes where institutions do not figure prominently. In other fields such as political economy, labor, industrial organization, international and macro, similar transformations have taken place.

But whereas the notion that institutions matter and are susceptible to analytical treatment has taken root, a second major theme in the institutional research agenda has not seen as much progress. This is the theme of how institutions arise and change over time. In North (1990) and Ostrom (1990) the importance of institutional change is already expressed in the titles of the books. North's title explicitly refers to 'institutional change' and Ostrom's to 'the evolution of institutions.' Despite this emphasis in the titles, even in these books the treatment of institutions as endogenous is much less developed than the theme of institutions as important exogenous constraints. In part this may be due to the fact that one theme naturally precedes the other, and also because institutional change is a much harder subject. Furthermore, theories of institutional change are a much harder sell to a readership that is accustomed to an equilibrium view of the world, where things can be controlled and predicted. Whereas the notion of institutions as exogenous constraints could be more or less easily grafted into existing theories with not too many incompatibilities with fundamental assumptions, institutional change seems to require greater departures from canon.

Aware of the shortcoming of our understanding of institutional change both North and Ostrom, as well as several other researchers, have sought to develop better and more explicit theories of how institutions emerge and change over time. These efforts are more forcefully expressed in North (2005) *Understanding the Process Economic Change* and Ostrom and Basurto (2011) *Crafting Analytical Tools to Study Institutional Change*. Both of these works focus centrally on institutional change and both explicitly take a Darwinian evolutionary approach together with some elements from the theory of complex adaptive systems.[1]

Although both of these authors were by then Nobel Prize laureates and the pieces dealt with what is widely considered a crucial issue, the reception was lukewarm at best and the impact on the scientific literature was much less pronounced than that of the earlier books. I do not have hard numbers to back these claims, but my personal experience talking to other researchers in the field is that North (2005) is hard to read and to make sense of, with its unfamiliar notions of non-ergodic worlds, beliefs, mental models and emphasis on cognition. If relaxing assumptions of the strength agents' rationality appears to be slowly gaining acceptance in mainstream Economics, the same cannot be said of out-of-equilibrium theories. Mainstream Economics and several other approaches in the social sciences have been built on a physical mechanics metaphor that assumes reductionism, linearity, Gaussian distributions, optimality and the ability to control and predict. This makes it very hard for those weaned in that tradition to accept evolutionary and complexity approaches, which are inherently non-reductionist, non-linear, full of fat-tailed distributions, where all the options cannot be specified and compared beforehand to yield the optimally choice. And most of all it is hard to get used to the notion of systems that never settle to a fix point but are constantly churning with incessant unpredictable novelty.

[1] For example, Ostrom and Basurto (2011) conclude their paper as follows: "Developing better tools to study the evolution of institutions is one important step we can take to reduce emphasis on institutional monocropping that currently dominates much of social science thinking as well as that of development agencies. We need multiple ways out of this trap. As academics, we can help by being willing to develop more complex theories for explaining the behavior of humans in widely divergent settings (Wilson, 2002; Schmid, 2004). We do not need to be complex, however, just for the sake of being complex, but we do need to get over our simplicity hang-ups. Obviously, our theories will always be simpler than the worlds we study, or we are trying to reproduce these worlds rather than a theory of these worlds. Given the complex, nested systems of the biophysical world, however, we need to develop a social science of complex, nested systems."

The position taken in this paper is that North (2005) and Ostrom and Basurto (2011) are on the right path for developing a theory of institutional change by focusing on the evolutionary process of variation, selection and replication through which institutions and beliefs change over time. The purpose of the paper is to propose a way to look at this evolutionary process using some analytical tools from the theory of complex adaptive systems, namely, fitness landscapes and coevolution. The hope is that the additional cost to the reader of understanding these possibly somewhat obscure methods will pay off in terms of new perspectives on some of the crucial concepts and ideas that North, Ostrom and others have been pursuing in the quest for a better theory of institutional change. Fitness landscapes provide a highly visual and intuitive way of perceiving some tricky concepts such as beliefs, ergodicity, adaptive efficiency, path dependence and even the process of economic development.

In this the purpose of the paper is similar to Mokyr (2014) which takes North's *Understanding the Process of Economic Change* and proceeds to "unpack these definitions and come up with a meaningful and useable definition of some Northean terms that could bear a bit of clarification (Mokyr, 2014: 151)." Importantly, Mokyr does this by looking and culture and beliefs through the perspective of cultural (Darwinian) evolution. The difference is that he illustrates these ideas by applying his framework to the British Industrial Revolution and I use fitness landscapes.

The paper is structured as follows. The next section (section 2) defines beliefs and how they relate to institutions. It also motivates the view that beliefs are a crucial determinant of whether countries develop and prosper or fail to do so. Section III then shows how beliefs can be understood as a set of instructions and thus treated as a bit string (of zeros and ones) that evolves on a combinatorically large fitness landscape. The shapes of these landscapes – simple, rugged or chaotic - mirror the nature of the problem of searching for fit design of beliefs, so the determinants of those shapes will be explored and reasons why evolutionary processes tend to yield rugged landscapes will be discussed. The key determinant will be the extent of interdependencies among the dimensions that make up a belief. But while rugged landscapes represent complicated problems that are hard to solve, they are not complex (in the sense of complex systems). It is only when we bring in the fact that a society's beliefs coevolve with those of other societies, as well as with technology and other systems, that landscapes are no longer static and start to dance. Dancing landscapes imply perpetual change and thus the need for constant adaptation. These concepts provide a model or a metaphor that give a fresh perspective to difficult ideas such as ergodicity and adaptive efficiency. The paper concludes by arguing that this approach to beliefs and institutional change suggests that development will probably remain an elusive goal for most countries, echoing Douglass North when he warned at the outset of his book that "the argument of this study suggests a sobering appraisal of the future of humans in the face of ubiquitous uncertainty of a non-ergodic world (North, 2005: 8)."

2. Beliefs, Institutions and Development

In order to analyze the process through which beliefs and institutions evolve using fitness landscapes it is first necessary to define and contextualize these concepts. This will be done by relating beliefs and institutions to the process of economic development and the question of why development is so elusive given the stock of our knowledge.

Cross-country data on GDP per capita shows that there is a small group of high-income countries and a large group of low-income countries, with few transitioning countries in between. In 2012 there were 26 countries with GDP per capita above US$25,000 – excluding oil producers and tax havens.[2] At the same time there were 143 countries with GDP per capita under US$15,000 and only 12 countries making the transition in the US$15,000 to US$25,000 range. Furthermore, most of the countries in the high-income group achieved this relative prominence in the 19th century, with few new countries joining the high-income group in the past century. Although economic theory presents many theoretical reasons why there should be a convergence of poorer countries towards higher incomes, such as diminishing returns and second-mover advantages, there is very little evidence of catch-up as of yet. On the contrary, both GDP per capita as well as more specific indicators of economic development indicate that most countries have not been making significant strides towards closing the gap. Most countries in this lower group seem to go through cycles of booms and busts where bad policies, missed oportunities and inefficiencies are endemic, suggesting that development is not just a matter of time (see

[2] Data in current US$ from World Bank national accounts data, and OECD National Accounts data files available at http://data.worldbank.org/indicator/NY.GDP.PCAP.CD.

Douglass North's epigraph to this paper) . Why is it that development is so elusive despite widespread new technologies and greater knowledge as well as the example from countires that have manged to make the transition?

The nature and causes of the wealth of nations has been a major issue in Economics at least since Adam Smith. Many different hypotheses about the fundamental determinant of long term growth and development have been suggest, including; tecnology (Solow, 1957); capital and savings (Harrod, 1939; Rostow, 1959); human capital (Romer, 1990, Glaeser et al. 2004); geography (Sachs and Warner, 1995); culture (Landes, 1998); factor endowments (Engerman and Sokoloff, 2000); institutions (North, 1990; Acemoglu and Robinson, 2012); among others. Currently it seems as if the instiutions-as-the-fundamental-determinant-of-long-term-growth hypothesis has come out on top.

But if weak institutions are in fact such a major impediment to long term growth and development, why is it that so many countries systematically fail to set up growth-enhancing institutions and maintain instead institutions that perpetuate inefficiencies and backwardness? North (2005: 3) argues that there are two answers to this question:

> Throughout history and in the present world economic growth has been episodic because either the players' intentions have not been societal well-being or the players' comprehension of the issues has been so imperfect that the consquences have deviated radically from intention.

The first of these answers is by far the most commonly invoked reason for backwardness and underdevelopment. Elites in these countries purposefully choose inefficient institutions due to the political risk inherent in putting into place different institutions. The new arrangements may empower opponents and bring about scenarios where the current elites would be worse off. Although compensating payments or Coasean bargaining could make the changes Pareto optimal, political transaction costs and the general difficulty in making credible commitments makes the current winners prefer a large slice of a smaller but more certain pie, than a uncertain slice of a larger pie. These arguments centered on political power and the inherent difficulty of transacting away inefficiencies are the basis of much of the political economy literature and in many cases are a sufficient explanation for the diversity of performance among nations.

However, though less popular, North's second reason for poor performance is also often important. By 'players' comprehension of the issues' he refers to the beliefs held by society about how the world works, which guide the choice of institutions, which in turn crucially contribute to determining economic performance. But if institutions can determine outcomes, how do countries choose which formal institutions to set in place? Most of the literature has addressed this issue by focusing on the relative power and interaction of the configuration of groups in a given society. The choice of insitutions is determined by the interplay as those groups with more power seek to put into place rules conducive to their own view of society, usually assigning most of the benefits to themselves and perpetuating their own political ascendency. But achieving the intended objectives is not usually that simple, for besides the intergroup struggle to prevail in the choice of institutions, there remains the serious difficulty in knowing which set of institutions will actually promote those outcomes. Before knowing what rules to strive for, a group has to have a notion of what set of institutions will produce the kind of outcomes they desire. I define this perception of cause and effect between institutions and outcomes as 'beliefs.'[3] This is similar to what Denzau and North (1994: 4) have called shared mental models: "The mental models are the internal representations that individual cognitive systems create to interpret the environment and the institutions are the external (to the mind) mechanisms individuals create to structure and order the environment." Mokyr (2014: 154) provides another, yet compatible, definition of beliefs, where culture is composed of beliefs, values and preferences. In this definition **beliefs** "contain statements of a positive or factual nature that pertain to the state of the world, both the physical and metaphysical environment and social relations. **Values** pertain to normative statements about society and social relations, and **preferences** are normative statements about individual matters such as consumption and personal affairs." Other related analyses of economic growth and development where beliefs figure prominently with more or less similar definitions are Grief (2006), McCloskey (2006, 2010), Mokyr (2009), Schofield (2006). In the next section we will see that fitness landscapes provide a very intuitive visualization of the metaphor often invoked in North (2005) of beliefs as scaffolding on which humans stand to erect institutions.

For the purposes of this paper a 'shared' or collective belief held by most of society emerges during periods of uncertainty and disruption when the previous set of beliefs has disspated due to poor fit to realized outcomes (Alston et al. 2016). The process through

[3] For a more detailed definition of beliefs along these lines see Alston et al. (2015).

which this happens will be analyzed below using fitness landscapes that graphically depict institutions, beliefs and outcomes and the punctuated process of change that emerges. It will be argued that although this evolutionary process is a powerful mechanism for searching fit designs in the case of biology, technology, culture and language, in the case of beliefs, institutions and development this 'algorithm' has not guided most countries to sustained development. While the persistence of sub-optimal institutions may be primarily due to vested interests that impede the changes and reforms that would induce growth and prosperity, the difficulty of revising and updating misguided beliefs is additionally responsible for the surprisingly small number of countries that have transition from underdevelopment to sustained growth in history. Thus even in those rare cases where a middle-income country reaches a situation where the fight for rents is kept in check and government policy is primarily interested in promoting the common good, there are still formidable impediments for achieving the transition.

3. Fitness Landscapes as Search Problems

Fitness landscapes, or adaptive landscapes, were first developed by Sewall Wright (1932) and have been widely used in evolutionary biology. This section describes what they are and how they can be applied to the case of beliefs.[4]

As defined in this paper beliefs are a mapping from institutions to outcomes. Those in power know what outcomes they desire but they do not know which institutions are most conducive to those outcomes. Therefore they have to have a view of how the world works that will tell them the cause and effect relation between each conceivable set of institutions and each set of outcomes. Suppose that institutions can be formulated as composed of N different dimensions. Acemoglu and Robinson (2012) have popularized the classification of institutions into two dimensions – economic and political - that can be either extractive or inclusive. But in principle there can be several other dimensions, such as, markets vs. state, secular vs. spiritual, progressive vs. conservative, centralized vs. decentralized, public

[4] While some approaches see fitness landscapes as a metaphor for how evolution works, others treat it more rigorously as a model. There is some debate in the literature about how useful they are for analytical purposes (for two sides of this debate see Pigliucci (2012) and Skipper and Dietrich (2012)). One objection is that we cannot really quantify fitness meaningfully, so the landscapes cannot be constructed for real cases. But note that this is not very different from utility theory in economics: the fact that we cannot measure utility or make interpersonal utility comparisons does not stop the theory from being a useful way to analyze choice and behavior.

regarding vs private regarding, role of the army, trust, race, gender, traditional vs. forward looking, etc. If there are *N* dimensions, and assuming for simplicity (as did Acemoglu and Robinson (2012)) that each dimension can take on two levels, then each set of institutions can be written as a bit-string of zeros and ones (which enables computation). If there are *N* dimensions, then there will be 2^N different possible institutions. Even if *N* is a relatively small number, this means that there is an inordinately large number of possible permutations of institutions for a country to choose from. If *N=15* then there are 32,768 alternatives. If one defines each dimension more finely, allowing for say *N*=30, the total number of alternatives would be over 1 billion. If each dimension is gradated into more than two states (zero/one or on/off) then the number of possibilities becomes unwieldly.[5]

In order to construct a fitness landscape take a given bit-string defining a specific set of institutions, say 0000000000, and plot it on a plane. Then plot adajcently all of its *N-1* one bit mutation neighbors, that is, all the adjacent set of institutions that are identical to it except for one dimension that is flipped from zero to one or vice-versa, such as 1000000000 and 0100000000, etc.[6] This plane containing all the possible combinations of institutions is known as the space of institutions.[7] For each point on the plane plot on the vertical axis its fittness. In biological applications fitness measures the entity's capacity to reproduce given the environment. In the case of institutions and beliefs it is a bit more subtle. It is tempting to think of the fitness of a set of institutions as the social welfare it contributes to generate, so that the fitness landscape can be seen as a welfare function. But more rigorously fitness refers to the capacity of that set of institutional arrangements to promote the expected outcomes, according to the beliefs, and thus to be kept in place, that is, its ability to perpetuate into the future. Another way to put this process of differential reproduction through which cultural entities evolve, according to Skyrms (2014: xiii) is that "successful

[5] This abundance of possibilities is what Ostrom and Basurto (2011: 324) refer to when they are analyzing the evolutionary process through which rules change and they state: "Given the logic of combinatorics, it is impossible for public officials, or beneficiaries, to conduct a complete analysis of the expected personal benefits, or broader performance of all the potential rule changes that could be made… A similar possibility also exists for biological systems – they evolve."

6 Ostrom and Basurto (2011: 329) follow a similar procedure to encode an action situation as a set of *N* rules which can take the value of 0 (no rule), *R* (required), *P* (permitted) or *F* (forbidden) and then examines the 'fit of rules to biophysical and community characteristics of a particular setting'.

[7] For a technical treatment of fitness landscapes see Kauffman (1995). The use of a plane to designate all the possible combinations is a simplification that is intuitive but not mathematically rigorous. A more formal treatment would use Boolean hypercubes (see Kauffman 1995).

strategies are communicated and imitated more often than unsuccessful ones." Figure 1 shows three different classes of fitness landscapes: simple or single-peaked, rugged and chaotic. Note that only a small portion of each landscape was drawn, as they can strech out towards infinity.[8]

[Figure 1 here]

What determines the shape of the fitness functions and what does the shape tell us about the problem that is being solved by the evolving entity? A peak on the landscape indicates a fit design of institutions, that is, a good way of organizing a society relative to other nearby arrangments. It ranks vertically which designs are better than others in terms of being replicated into the future. The three landscapes in Figure 1 show different possibilities. The first landscape, sometimes called a Mount Fuji landscape, has a single-peak and smooth contours. It represents a situation where there is a single global optimum. Because that single solution is easily perceived – it can be found by a simple hill-climbing algorithm – it represents a simple problem that can be solved.[9]

The middle figure is a rugged landscape that has several local optima. This is a situation where there are many potential solutions and it is not always clear which one is the best. Because the landscape can be very large and unobservable it is not possible to simply scan the surface and choose the best solution. It is also not possible to try out every possibility to determine which point is best. Instead, some kind of search method will have to be used to try to reach a good solution without getting stuck in inferior local optima, which a simple hill-climbing algortihm would do. It turns out that evolution, that is, an algorithm based on variation, selection and replication, is particularly well tuned to find fit design in rugged landscapes. Selection climbs to higher ground while variation (through mutation and recombination) allows for exploration that evades inferior local peaks.

To visualize how evolution would look like on a fitness landscape imagine a group of dots distributed over a portion of the landscape, representing a population of organisms. They are grouped closely because they are the same species, but are not all on the same

[8] A fitness landscape for a human genome, which is composed of approximately 20,000 genes, can have more than $2^{20,000}$ combinations, which is more than the number of atoms in the known universe (10^{80} atoms).

[9] Most of standard economic theory takes place in a Mount Fuji environment as there is a single fixed point that provides the best solution. This is not because the object being studied – the economy -actually has these properties, but due to an epistemological preference for equilibrium solutions. The single-peak is obtained through strong assumptions of convexity, monotonicity and continuity.

point because each individual has slightly differrent traits, due to recombination and mutation. Those whose traits are more adapted to the environment will be in a higher point on the landscape and will reproduce at a higher rate, passing on those traits to the next generation. This differential reproduction is the process of evolution. Visually the higher points will spawn other point that appear to climb the peaks of the landscape and the lower points will disappear. In a Mount Fuji landscape this leads to the species at or near the single peak. On rugged landscapes the outcome is sensitive to intial conditions, path dependent and will most probably not reach the global optimum, though local optima will likely be reached.

The third landscape is chaotic or random. Contrary to the other two, the fitness of a given set of institutions is in no way related to the fitness of neighboring institutuions, leading to sharp spikes and valleys. These landscapes represent problems that are practically impossible to solve, except by chance, as being on one point does not give you any information of the fitness of any other point. In such an environment evolution would not be a good search algorithm.

Figure 1 shows that even if one is able to find the optimal solution in each of the three types of landscapes, the quality of that solution is higher for Mount Fuji landscapes than for the rugged landscapes, which in turn is higher than that for random landscapes, that is, the scale of the figures is drawn to be comparable. In other words, simple problems allow for good solutions, while even the best solution for hard problems are not so satisfactory. Why would that be so? What is it that holds down the peaks in the rugged and random landscapes? The answer to this question turns out to be the same thing that determines the shapes of the landscapes: the number of links between the dimensions of the institutions. Recall that an institution is composed of N dimensions arrayed as a bit string of *0*s and *1*s. Starting from a random bit string, such as that in Figure 2, how would one go about changing it in order to reach the best possible configuration? If the contribution of each dimension to the institution’s total fitness is independent from the contribution of every other dimension, that is, the dimensions are orthogonal, then one could flip one of the dimensions from zero to one (or vice versa) and keep the change if the fitness increases or switch back otherwise. Doing this for each dimension would take one to the optimal set of institutions. When the dimensions are orthogonal in this way, the corresponding landscape is a Mount Fuji, with a single best solution that yields high levels of fitness and is easy to discover.

[Figure 2 here]

But in most situations it is natural for some dimensions to be interrelated to some of the other dimensions. The political dimension of a set of institutions is never perfectly orthogonal to the economic dimension or the racial dimension, for example. In genetics these interrelations are called epistatic couplings, but other useful labels are 'externalities across dimensions' or 'internal constraints'. When there are links between the dimensions of an institution the exercise of flipping one dimension at a time and checking for improvements won't work, for each additional flip can undo the previous improvements. The search problem is much harder now. Furthermore, because there are more constraints the solutions - even the optimal solution - will generally be lower than the case where there are no constraints. What differentiates the rugged landscape from the random landscape is the number of links. If K is the number of links between each dimension, simple landscapes have $K=0$, rugged landscapes have positive yet low K, and random landscapes have complete links, that is, $K=N-1$. In Figure 2, for example, $K=2$.[10]

We will see in the following sections that there are good reasons to expect many of the problems we study in social science to be rugged landscape problems, as is the case in biology. That is, these are situations that are not in the extremes of simplicity and chaos, but instead in the interesting complex region at the edge of chaos. I have been careful to call the rugged landscapes 'complicated' rather than 'complex', even though their rugged nature is caused by connectedness, interdependencies and diversity, three of the four basic key components of a complex system. But the fourth component, **adaptation** or **learning**, is still missing, for the landscapes have been described as fixed or static, as that they do not react to what the evolving agent or anyone else is doing. Below I will introduce additional interdependencies into the system that will cause the landscapes to dance, that is, adapt to what others are doing, and then the system will become truly complex. First, however, let us see how we can use the concepts of landscapes to understand the notion of 'beliefs' and how they evolve.

[10] For further discussion on the determinants and the implication of the shape of fitness landscapes see Kauffman (1995), Beinhocker (2006: chap. 9) and Page (2011: chap.3)

4. Beliefs as Information Evolving on a Fitness Landscape

Perhaps the most controversial aspect of applying an evolutionary approach to any specific area is determining what is the unit of selection on which evolution operates. In biology the original Darwinian theory took the unit of selection to be the individual or the organism, that is, the phenotype. This is the way that most people tend to think about evolution. An animal which, through mutation or breeding, becomes better/faster/stronger than the others with which it competes for resources, will have greater reproductive success leading to more of its kind. But this view was eventually challanged by the notion that evolution acts instead on the genotype, that is, on the set (or even a subset) of genes in an organism (Dawkins, 1976). In this view the individual is just the vehicle for the change and it is instead the genes that encode the information for replicating the phenotype that should be the unit of analysis, as they are the entity of which copies are made and passed along. Other have argued that, in fact, any entity that goes through heritable variation and selection can be subject to an evolutionary process, so that the unit of selection can be the genes, chromosomes, genomes, individuals, groups, demes, populations, species or even ecosystems. Evolutionary biology has been through fierce disputes over which is the proper unit of selection (known as the 'units of selection debates') and it is probably the case that different units of selection are important for analyzing different questions.

One of the most appealling applications of evolutionary logic to epsitemic substrates has been that of technology and technological change. In evolutionary accounts of the process of variation, selection and amplification of technology it is tempting to think of the artifacts that arise from the efforts at invention and innovation as the unit of selection. However, Mokyr (2000: 6-7) proposes instead that the proper unit of analysis is the technique, that is, "the set of instructions, much like the if-then algorithms of a computer program, that tells someone how to produce, that is how to manipulate the forces of nature in the interest of material well-being of people". In a similar manner, when it comes to analyzing economic development of nations one may be tempted to consider that it is a country that goes through the process of variation, selection and replication. Alternatively one might see the country's institutions as the unit of selection. However, I follow the logic of the variants of the biological and the technological analyses that put the 'set of instructions' as the unit on which evolution operates. The 'set of instruction' on what institutions to put in place in order to produce the sought after material well-being is precisely what this paper has defined as beliefs. Beliefs are thus akin to the genotype in biology and to the technique in technological change, while the institutions are the vehicle, that is the rendition of those beliefs.[11]

In order to explore how beliefs emerge and change let us return to the landscape representations and suppose a situation where a nation has just undergone a major disruption due to poor performance that ousts the ruling elite creating a window of opportunity for a change in the configuration of power and in institutions. Eventually a configuration of groups will rise to power forming an elite and will be confronted with the choice of which formal institutions to put in place. In order to do this they would like to examine the landscape that shows the fitness of each possible combination of institutions and choose the set that corresponds to highest point. But unfortunately the landscape is not observable. The only way to know the landscape is to test each combination of institutions to ascertian what outcomes would result. But that is too costly and difficult, even for a small subset of the entire set of possible institutional arrangements. Even relying on historical experience or that of other countries would not work as the landscape is specific to a given circumstance and time.

Nevertheless a decision has to be made: some set of formal institutions has to be put in place. To make that choice a society necessarilly has to infer what they think is the shape and location of the current instituional fitness landscape. In order to do so they will use their perception of how they think the world works, which will have been shaped by uncountable factors, such as historical experience, recent traumas, national charcater, culture, religion, fads, global trends, etc. This perception of how the world works is what I have been calling beliefs. Note that by this definition beliefs can also be represented by a landscape that maps from institutions to expected fitness. In fact, beliefs can be represented simultaneously on the same graph as the actual outcome surface. If the perception of how the world works actually fits reality closely, then the two landscapes will be very similar to each other. But

[11] Ostrom and Basurto (2011: 325) take a similar strategy with 'rules': "A rule configuration is parallel in function to a genotype in that rules like genotypes are mechanisms that transmit information about how to produce something (a protein of a particular organ or an action situation in a particular environment. In other words, rules are a set of instructions of how to produce the expressed situation or the structure of the relationship among individuals that is also affected by the biophysical world and the kind of community or culture in which an action situation is located."

in general this perception will be distorted by all sorts of uncertainties, unknowables and biases so that the two landscapes can diverge markedly.

Figure 3 shows a stylized example where the beliefs and actual outcomes diverge. In this example beliefs are always above what will really materialize if those institutions are put into practice, that is, those beliefs are overly optimistic.[12] Given the beliefs in the upper landscape in Figure 3, the best possible outcomes that can be expected would reach a fitness at point *A*, measured vertically. This could be achieved by putting into place institutions corresponding to point *C* in the design space of institutions. Yet in making this choice this society is setting itself up for disappointment because rather than outcomes with fitness at height *A* what will actually emerge are outcomes corresponding to fitness level *B*. If the effect of institutions were readily observable, this society would realize that the best possible outcomes which they could aspire at this point in time correspond to point *E* which can be reached through institutions at point *F*. Given their beliefs, however, they see no point in trying institutions at *F*, which they erroneously expect to yield outcomes *D*, lower than those at *A* which they think that they can achieve.

[Figure 3 here]

Note that societies never really see the entire landscape of actual outcomes. All they see are the outcomes of those institutions that they actually try out. So in Figure 3 society sees only distance *AB*. This distance is a measure of the selective pressure in this evolutionary process. The greater the disconnect between beliefs and outcomes, the greater the disappointment and frustration of that society with the institutions they have adopted. This disappointment puts pressure to revise beliefs and can lead to incremental or punctuated adjustments. North (2005: 79) stresses that “an immense amount of economic change has been the unanticipated result of institutional change that reflected a significant gap between intentions and outcomes as a result of ‘faulty beliefs.’”

But feedback in these systems is slow and societies might hang on to their beliefs for a long time before they realize that things are just not going to turn out as expected. Cognitive dissonance often leads societies to conjure up justifications for why things haven’t yet turned out as expected given their beliefs rather than admitting those beliefs

[12] Note that this portrayal of beliefs as landscapes gives a visual counterpart to North’s literary image of beliefs as scaffolding: societies are standing on their beliefs and that determines where they will build their institutions.

were wrong. For these reasons the poor outcomes might persist for a long time, until an event takes place which triggers a sudden demand for change by large parts of society. Such events are windows of opportunity, which might take the form of a revolution or a coup, but also myriad other forms such as an electoral upset, a new constitution or even something more subtle. Eventually new beliefs will emerge leading to new institutions. Although it is tempting to see this process as leading to continual improvement over time, there is no reason to necessarily expect progress or development. On the contrary, the historical experience has been that most countries have continuously gone through this evolutionary process jumping from set of dysfunctional beliefs to another.

Whereas there is also no assurance of progress from evolution on substrates, such as biology, technology and language, in these cases outcomes have seemed substantially more adaptive, complex and functional than in the case of beliefs and institutions. Why would it be that evolution seems to be so slow for beliefs when they all face rugged landscapes. An important distinction is that these other areas try many experiments at a time, in a sense taking a shot gun approach to finding fit design on the landscape. A gene uses an entire population to try many different experiments at the same time. A society, on the other hand, can have only one shared belief at time and trial and error of institutions is thus a slow process. Furthermore, feedback of what works and what doesn't is slow and unprecise, so that nations can insist for long periods on strategies that will not work. It is true that unlike biological evolution where the variation is blind (mutation, crossover and other processes), change in epistemological evolution has intentionality and direction. Nevertheless there are still errors, surprises and accidents that create unintended variation. In making this point for the case of technological change Mokyr (1993: pp. 286) states that "there is a large accidental component in long-term change, and the record suggests that serendipity, opportunism, and the 'King Saul effect' (which occurs when the search for the solution of a specific problem inadvertently leads to a totally new opportunity set) plays an important role." To see this point simply consider how frequently institutional development programs go array, no matter how many of the best local minds and international consultants are involved and even when rent-seeking and politics do not get in the way.

Up to this point in the paper fitness landscapes have been taken as being static, thus representing a fixed problem. Even with this simplification the task of searching for fit design can be very difficult, especially, it has been argued, for the case of societies' beliefs on how the world works. Evolution turns out to be a powerful search strategy in rugged landscapes, but the nature of the problem is such that there is no presumption of progress or optimality. But the full picture is even starker than this. In the next section the landscapes will no longer be analyzed in isolation, but will be interrelated to the landscapes of other coevolving entities. This will make the landscapes dance so that the problem is no longer fixed but continuously churning.

5. Coevolution and Dancing Landscapes

You realize that for each species, its landscape consists almost entirely of other species, all of them busy evolving right back. That's co-evolution. We are all each other's fitness landscapes. Stuart Brand (2012: 125)

The height of an evolving entity on a fitness landscape is determined by how adapted its traits are to the environment in which it functions. We have been taking that environment as fixed and given. But every entity is nested in a network of relationships with other entities, who might be predators, prey, competitors, cooperators, symbionts, parasites, etc. Though these terms are most commonly used to refer to biological and ecological relationships, they are apt also for other evolving substrates such as beliefs, institutions and culture. These connected entities are themselves constantly evolving. In addition every evolving entity has its capacity to reproduce affected by changes in technologies, institutions, language, and even climate and geography. This means that the fitness of a specific configuration of traits for a given entity will change when any of these other linked evolving entities change themselves. That is, the landscape is endogenous and will not always remain fixed. Instead the landscape will dance, as what was a fit trait in a given environment becomes no longer such a good strategy for survival and reproduction. In the same manner, traits that were not well adapted to a given set of challenges might all the sudden become very advantageous when changes take place among competitors or cooperators. A spider's ability to make sticky web will start to contribute less to the spider's fitness if the fly evolves wings that do not adhere to the web, for example. And with this collapse in the spider's landscape there will be evolutionary pressure for it to evolve other traits that enhance its capacity for survival and reproduction, say an even stickier web, in turn making the fly's landscape recede. The result of this dancing is that what was already a difficult problem, finding fit design, becomes even harder. It leads to what is known as the Red Queen's Race, that is, a race in which "it takes all the running you can do, to keep in the same place," as noted the Red Queen in Lewis Carroll's *Through the Looking Glass* (1871).[13]

While the shape of the landscape is determined by the number of (*K*) epistatic links between the (*N)* dimensions of the evolving entity, the dancing of the landscapes is determined by the (*C*) external couplings between one entity's dimensions or traits to that of (*S*) other entities.[14] Figure 4 illustrates the simultaneous interrelatedness among the dimensions of a country's beliefs, and between those dimensions and the beliefs of other cooperating or competing countries, as well as with the dimensions of evolving technology (other entities could be added). Note that each evolving entity has internal epistatic links and external couplings. So now the system is truly complex, for in addition to connectivity, interrelatedness, and diversity, we now have adaptation/learning, as each entity's evolution ripples through the system changing the nature of the search problems faced by other entities in a tectonic landscape dance.

[Figure 4 here]

And because it is a complex adaptive system another property applies: complexity happens not at the extremes, but at the edge of chaos. This is so for both the internal epistatic couplings and for the external couplings among entities. For the case of the external links, note that if there are no couplings the landscape of each entity is fixed and doesn't dance at all. If the system is too highly interconnected, with most traits of each entity linked to most traits of all other entities, the landscapes thrust up and down too violently to sustain order and functionality. It is only in a very restricted range of external interconnections that the landscapes will move at rates where coevolution will yield interesting complex structures and relationships.

Regarding the internal links, one might expect evolution to select for configurations of coevolving entities that evolve on Mount Fuji landscapes, for as we have seen in Figure

[13] The relationship can also lead to an equilibrium of sorts where each entity has reached evolutionarily stable strategies where there will be no pressure for change (see Maynard Smith, 1982). This will lead to periods of stasis and fixed landscapes, but this can be knocked over by shocks and perturbations.

[14] In terms of the more rigorous models on which fitness landscapes are built, the case of fixed landscapes corresponds to Kauffman's (1995) *NK* model and the case of dancing landscapes corresponds to the *NK(C)* model.

1, they have higher average fitness and should therefore reproduce at a higher rate. But this result changes when we consider coupled dancing landscapes. To see this visualize a species on a single-peaked landscape. Evolution will start pushing this specie's design towards the peak. But before the population reaches the peak, the mountain will subside and pop up unpredictably somewhere else in the landscape. Evolution will now push the population towards that new design, but again, before the higher ground is reached the landscape dances again sending the single peak somewhere else. All the while the population is on low levels of the landscape and thus accruing low fitness with severe consequences on reproduction and survival. A similar result arises in the case of chaotic dancing landscapes. The population will find it impossible to evolve towards fit design for the needles and holes that make up the landscape are moving up and down randomly making average fitness very low. In the case of rugged landscapes (low but positive K) the population actually has a chance. The peak towards which a population is evolving might subside, but because of the rugged nature of the terrain, one or more other peaks will pop up close enough nearby to allow it to be reached. Even though the peaks are lower than the Mount Fuji peak, a mid-size peak that can be reached is better than a mountain that is always too far away. That is, on rugged landscapes the average fitness of the entities in a co-evolutionary relationship will be higher than that the average fitness of relationships with too few or too many internal and external couplings. If this is the case, then evolution should select for co-evolutionary relationships that have precisely these characteristics.

This result, which as far as I can trace is due to Kauffman (1995), is profound and bears repeating: coevolving relationships are themselves selected by an evolutionary process. It is in a sense meta-evolution. It implies that interesting, complex, functional co-evolving relationships – among genes, beliefs, institutions, technology, language, culture and others, and therefore also among their vehicles, species, societies and countries– will emerge in very restricted circumstances where interdependencies, connectivity, diversity and adaptation are just right. In "Where Good Ideas Come From: The Natural History of Innovation" Steven Johnson has identified these circumstance to come together in 'liquid networks' such as coral reefs, tropical jungles, coffee-houses, online and large cities. Our interest here, however, is in economic development and prosperity, so in the next section we will see how the theory of coevolving landscapes can help us better understand these issues.

6. Development as an Emergent Property of a Complex Adaptive System

> *…the very efforts of humans to render their environment intelligible result in continual alterations in that environment and therefore new challenges to understanding that environment. (North, 2005: 5)*

The key problem that both North (2005) and Ostrom and Basurto (2011) seek to explicitly deal with is the fact that the world is continuously producing novelty and change, so that beliefs and institutions that humans adopted in the past to reduce uncertainty will not necessarily work in the new reality. North (2005: chap. 2) calls this ever-changing reality a 'non-ergodic world'. Ergodicity is the probability that any state will recur, so that a non-ergodic stochastic process is one where the averages calculated from past observations can be persistently different from the averages of future outcomes. Ergodicity is a fundamental assumption of Economic theory.[15] North's (2015) critique is that agents with imperfect perception in a non-ergodic world – which is the human condition – will face systematic relationships that change over time in unpredictable ways so that new fundamentally different uncertainties will constantly arise, creating the need for new beliefs and institutions to reduce those uncertainties. But because perceptions are imperfect and the future will have a different and unknown underlying structure (it is non-ergodic), finding these beliefs and institutions is a remarkable difficult process that different societies will address in very different ways, not the least because of the path dependent nature of the process through which societies go about this.

The view of the complex nature of economic development explained in the previous paragraph is somewhat difficult to assimilate, even though it does acutely capture the essence of the problem. Not only are there terms that are unfamiliar to most, such as ergodicity, but the argument is very abstract and different from the usual framing of problems as an optimization subject to constraints given fixed preferences and a known and finite set of options or strategies. Let us see therefore if the concepts of fitness landscapes and coevolution can provide some insights for understanding this view of development. Agents' choice variables are institutions and these are crucially affected by agents' beliefs

[15] Davidson (1991) is a good source on the ergodicity assumption in Economic theory. North (2005: 19) notes that Paul Samuelson saw the ergodic hypothesis as "essential for a scientific economics".

of how the world works. We have already seen in the previous section how the space of institutions can be unmanageably large and how beliefs can be portrayed as a rugged fitness landscape that often deviate from the relationships that actually exist in the world. We have seen that these deviations act as selective pressures to abandon unfit beliefs – in a possibly slow and punctuated process – and evolve new solutions. But we have seen as well that the nature of the search problem is also constantly changing, for the landscapes are coupled and can therefore change unpredictably over time. This notion of dancing landscapes provides a very visual and intuitive notion of a non-ergodic world. Beliefs that once fit well to how things work can start to fail once the landscape moves. Furthermore, the way the landscape moved in the past provides little information on how it might move in the future so learning and adaptation will never be complete, rather must be a never-ending process.

But one might object that the picture that emerges from this analysis is too bleak and does not accord with what is observed across nations and throughout history. Although it is true that things seem to change over time with occasional innovations and surprises, there is still a remarkable amount of stability and predictability. The set of rich developed countries has been basically the same for nearly a century with few new members joining in and almost none falling out. Countries do have booms and bust, good times and bad, but it all seems to fluctuate around a very stable core. So how can we square the view of a non-ergodic world where landscapes continuously dance in unpredictable and uncontrollable ways with this notion of a very balanced world of enduring relationships? What follows is a conjecture of how these two views can be incorporated into the same framework.

Casual observation and much scholarly analysis indicate that the countries that have achieved higher levels of development and prosperity have institutions with very similar characteristics. Even though the details and flavors of the institutions vary considerably, some characteristics seem to be very prevalent among those countries that made it and absent in those that didn't, such as openness, free entry and exit, democratic governance, rule of law, competence criteria, etc. This suggests that some parts of the space of institutions will systematically have higher fitness. For an individual country this would be the case if those regions in its space of institutions that encode for institutions with variations of the above mentioned characteristics had on average higher fitness than other regions without those characteristics. Figure 5 shows a hypothetical fitness landscape of a country's beliefs and institutions. This landscape is not drawn as smooth as the previous ones but note that different regions have loosely correlated fitness levels, indicating that small changes in institutions have some impact, but that there is a common element where similar institutions have fitness at the same general level (e.g. high, medium or low).[16] In the figure there is a region where all combinations of institutions have relatively high fitness, with some variation among them. History suggests that these characteristics might be related to openness, rule of law, democratic governance, etc. Some other regions have pockets of low fitness, again with some variation among them. Again, history suggests that these are far more prevalent, covering a larger area of the space, and are related to coercive, status-based extractive institutions. This landscape is coupled to other landscapes so it will dance. But the conjecture here is that the dancing changes the fitness locally, making the institutions in the high/low fitness region become higher or lower than others in the same region, but maintaining the same general structure where some regions have fitter institutions than others. Metaphorically it's as if earthquakes in the Andes change which are the highest peaks, but all terrain in the mountain range is still higher than the plateaus to the east. As with geology, it is possible that in the very long run even the more stable core structure might change with some entirely new set of beliefs and institutions becoming more conducive to development and prosperity.

[Figure 5 here]

In addition, there is good reason to expect that the landscapes for different countries will also be similarly shaped in their core structure. This is so because all developed countries have reached that position through beliefs and institutions that have variations on the same general characteristics, i.e. openness and inclusion. But again, this does not mean that the landscapes are identical or even similar, as each country has its own history, shared experiences and physical constraints. North (2005: 66) notes that "different experiences of societies through time will produce different perceptions of the way the world works and therefore require different institutions to provide the same incentives." Furthermore, path dependence implies that even countries that started from very similar initial conditions can end up taking very different trajectories. This becomes readily apparent in Figure 5 by

[16] In other words, the landscape is not isotropic, that is, not invariant with respect to region.

placing two sets of beliefs next to each other and letting them evolve. The paths will very quickly diverge, especially if the landscape is dancing.

The image I am proposing here is therefore one where there are some common elements in different countries' fitness landscapes, but at the same time there is substantial dancing of the landscapes (or non-ergodicity) that uncertainties make reaching the higher regions and staying there, a very hard search problem. One might question why it is that countries, no matter where they are on the landscape, cannot simply 'jump' straight to the regions of high fitness. After all, unlike biological evolution institutions and even beliefs are not blind. There is intentionality and choice. The more our knowledge and experience grows the greater should be our understanding of how thing work and it should be easier establish better institutions. But just like a fish cannot evolve into a giraffe in few steps (it would take a very long process) there are constraints on the institutions that a country can choose. North (2005: 77) puts great emphasis on this "difficulty of fundamentally altering paths":

> It is more than simply that the organizations brought into existence by the existing institutional matrix owe their survival and well-being to that matrix and therefore attempt to prevent changes that would adversely affect their well-being. It is also that the belief system underlying the institutional matrix will deter radical change. (North2005: 77)

Because of path dependence beliefs and institutions are constrained to move not by leaps and bounds but rather by moving into the *adjacent possible*. This is a concept created Kauffman (2000) inspired by Jane Jacobs (1961) and popularized by Johnson (2010), which refers to the fact that because innovation is essentially a recombination of existing elements – be they technologies, ideas, new species, new materials and even beliefs and institutions – at any point in time there is a limited set of possible new elements or innovations that can be created. The adjacent possible is thus composed of all the first-order combinations that can be made with the elements you currently have. The more elements you have the greater will be the adjacent possible, and it grows with each innovation. The adjacent possible is thus "a kind of shadow future, hovering on the edges of the present state of things, a map of all the ways in which the present can reinvent itself" (Johnson, 2010: 31).[17] It "captures both the limits and the creative potential of change and innovation." Because of path dependence, therefore, societies face a limited adjacent

[17] For a related notion relevant for Growth Theory see Martin Weitzman's (1998) *Recombinant Growth*.

possible in which to advance. This is why best practices, institutional transplants and advice from international experts have not had a good track record in promoting development.

What is important for development to materialize and persist according to North (2005) is not so much the specific content of the beliefs, as there can be many different functional recipes, but rather a characteristic that he calls *adaptive efficiency*: "the ability of some societies to flexibly adjust in the face of shocks and evolve institutions that effectively deal with an altered 'reality'" (North, 2005: 6). Given that the landscape will dance, it is not enough to simply hit on a good solution for today's problems. What is required are institutions that readily adapt to shocks and the novel situations and uncertainties that constantly arise. Adaptive efficiency allows creative destruction to take place by facilitating commitment, intertemporal political transactions and other means to compensate blocking losers or overcome resistance to change. In this sense it may be that the key dimension of a country's institutions and beliefs (or epistatic interactions among dimensions) is that which gives it the capacity to adapt to change and move from places where the landscape has collapsed to others where it is surging.

7. Conclusions

One of the challenges in complexity theory is how to measure the complexity of an object, as there are too many measures and none captures all aspects. One of these measures of complexity, in computation theory, is called *depth* and captures the notion that complex objects require a lot of time (or memory) to produce. According to Mayfield (2013: ??) it provides a formal measure of the internal relationships that exist within a computational object.

> "The more complicated the internal relationships, the more steps it takes to compute an output having those complications from input that doesn't have them. This property is known as 'slow growth,' and it means that deep objects can be created quickly only from deep inputs. Creation of complicated (deep) objects from simple (shallow) inputs requires a lengthy computation. There are no shortcuts to this rule." Mayfield (2013: ??)

This concept is extended from computation to other areas in a straightforward way. The process of development can be thought of as a process of computation, as the search for fit design described above is essentially a process of transforming information as inputs, different beliefs that are tried out in practice, into information as outputs, the knowledge of which strategies are better than others in those circumstances. The notion of depth suggests that reaching complex beliefs and institutions that are able to react quickly to new situations is a slow and time consuming process. Note that much of the time required to compute complex action and complex structures can sometimes be spent in advance. Mayfield (2011: ??) contrasts the encounter of a lion and a gazelle with an encounter between a lion and a prepared human. “The human cannot outrun the lion, but may carry a weapon. The gun or spear took much time to create, but that time was spent before, not after, a lion is encountered. The complexity inherent in the weapon is created and banked for future need.” Similarly in development much of the knowledge acquired is embodied in technology and in books, blueprints and recipes, but the crucial and element is the tacit collective know-how that is held diffusely by networks of people and organizations (Hidlago and Hausmann, 2009; Haumann and Hildalgo, 2011). These complex networks are emergent phenomena that require appropriate beliefs and institutions to emerge.

This treatment of complex structures as taking long processes to create seems to be exactly what Douglass North refers to in his Nobel Prize Lecture (North,1993) when he states that “successful political/economic systems have evolved flexible institutional structures that can survive the shocks and changes that are a part of successful evolution. But these systems have been a product of long gestation. We do not know how to create adaptive efficiency in the short run”. And similarly in North (2005: 78) when after describing adaptive efficiency he states that “it is important to understand that we do not know how to create these conditions in a short period of time … in the Western world it has been an evolutionary product of centuries of institutional change.”

Note that deep complexity will only arise when the challenges faced by the agents is large. In simple environments deep structures are not necessary and the cost of evolving them is not worthwhile. But when it is a case of agents competing with each other, that is coevolution, there is always a need to deal with novel challenges and in this environment deep structures determine survival. This implies that coevolution should lead to the increase in the depth of complexity observed over time.

References

Acemoglu, Daron, and James Robinson. 2006. *Economic Origins of Dictatorship and Democracy*. New York: Cambridge University Press.

Acemoglu, Daron, and James Robinson. 2012. *Why Nations Fail; The Origins of Power, Prosperity and Poverty*. New York: Random House.

Alston, Lee J., Marcus A. Melo, Bernardo Mueller and Carlos E. Pereira. 2015. *Beliefs, Leadership and Critical Transitions: Brazil 1964-2014*. Princeton University Press.

Beinhocker, Eric D. 2006. *The Origin of Wealth: The Radical Remaking of Economics and What it Means for Business and Society.* Harvard Business School Press.

Brand, Stuart. 2012. "Fitness Landscapes." Editor John Brockman, *What is Your Favorite Deep, Elegant, or Beautiful Explanation?* New York: Harper Perennial: 122-125.

Carroll, Lewis. 1871. *Through the Looking-Glass and What Alice Found There*. Macmillan & Co.

Davidson, Paul. 1991. "Is Probability Theory Relevant for Uncertainty? A Post-Keynesian Perspective." *Journal of Economic Perspectives*. 5(1), Winter: 129-143.

Dawkins, Richard. 1976. *The Selfish Gene*. New York: Oxford University Press.

Denzau, Arthur, and Douglass C. North. 1994. "Shared Mental Models: Ideologies and Institutions." *Kyklos* 47: 3–31.

Engerman, Stanley and Kenneth Sokoloff. 2000. "Institutions, Factor Endowments, and Paths of Development in the New World", *Journal of Economic Perspectives*, 14(3): 217-32.

Glaeser, Edward L., Rafael La Porta, Florencio Lopez-de-Silanes and Andrei Shleifer (2004). Do Institutions Cause Growth? *Journal of Economic Growth*, v9, n3, September 2004, 271-303.

Grief, Avner. 2006. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press.

Harrod, Roy F. 1939. "An Essay in Dynamic Theory". *The Economic Journal* **49** (193): 14–33.

Hausmann, Ricardo, and Cesar A Hidalgo. 2011 "The Network Structure of Economic Output" *Journal of Economic Growth*. 16:309-342.

Hidalgo, Cesar, and Ricardo Hausmann. 2009. "The Building Blocks of Economic Complexity." *Proceedings of the National Academy of Sciences of the United States*, Vol. 106, No. 26: 10570-10575.

Johnson, Steven. 2010. *Where Good Ideas Come From: A Natural History of Innovation*. New York, Riverhead Books.

Jones, Charles I., and Paul M. Romer. 2010. "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital." *American Economic Journal: Macroeconomics*, 2(1): 224-45.

Kauffman, Stewart. 1995. *At Home in the Universe*. New York: Oxford Univ. Press.

Landes, David S. 1998. *The Wealth and Poverty of Nations: Why Some are so Rich and Some so Poor*, W.W. Norton & Company.

Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge University Press.

Mokyr, Joel. 2000. "Natural History and Economic History: Is Technological Change an Evolutionary Process?" Northwestern University, draft lecture http://faculty.wcas.northwestern.edu/~jmokyr/papers.html.

Mokyr, Joel. 2014. "Culture, Institutions, and Modern Growth." In Editors Sebastian Galiani and Itai Sened, *Institutions, Property Right, and Economic Growth: The Legacy of Douglass North*. Cambridge University Press: 151-191.

North, Douglass C. 1990. *Institutions, Institutional Change and Economic Performance*. New York: Cambridge University Press.

———. 2005. *Understanding the Process of Economic Change*. Princeton, NJ: Princeton University Press.

North, Douglass C., John Joseph Wallis, and Barry R. Weingast. 2009. *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History*. New York: Cambridge University Press.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.

Ostrom, Elinor and Xavier Basurto. 2011. "Crafting Tools to Study Institutional Change." *Journal of Institutional Economics*. Vol. 7, Issue 3: 317-343.

Romer, Paul M. 1990. "Endogenous Technological Change," *Journal of Political Economy*, October, 98(5), S71–S102.

Rostow, Walt W. 1959. "The Stages of Economic Growth." *The Economic History Review*, Vol.12, No. 1: 1-16.

Page, Scott E. 2011. *Diversity and Complexity.* Princeton University Press.

Pigliucci, Massimo. 2012. "Landscapes, Surfaces, and Morphospaces? What are they Good For?" in eds. Erik Svensson and Ryan Calsbeek. *The Adaptive Landscape in Evolutionary Biology.* Oxford University Press.

Sachs, Jeffrey and Andrew M. Warner, 1995. "Natural Resource Abundance and Economic Growth," NBER Working Paper No. 5398.

Skipper, Robert A. and Michael R. Dietrich. 2012. "Sewall Wright's Adaptive Landscape? Philosophical Reflections on Heuristic Value." in eds. Erik Svensson and Ryan Calsbeek. *The Adaptive Landscape in Evolutionary Biology.* Oxford University Press.

Skyrms, Brian. 2014. *Evolution of the Social Contract*. Second edition. Cambridge University Press.

Schofield, Norman. 2006. *Architects of Political Change: Constitutional Quandries and Social Choice Theory*. New York: Cambridge University Press.

Solow, Robert M. 1957. "Technical Change and the Aggregate Production Function". *Review of Economics and Statistics* The MIT Press **39** (3): 312–320.

Weaver, Russell and Jason Knight. 2014. “Evolutionary Mismatch as a General Framework for Land Use Policy and Politics” *Land*. 3, 504-523.

Weitzman, Martin. 1998. “Recombinant Growth” *Quarterly Journal of Economics*, Vol. CXIII, May, Issue 2: 331-360.

Wright, Sewall. 1932. “The Roles of Mutation, Inbreeding, Crossbreeding, and Selection in Evolution”. *Proceedings of the Sixth International Congress on Genetics*. pp. 355–366.

Figure 1 – Simple, Rugged and Chaotic Landscapes

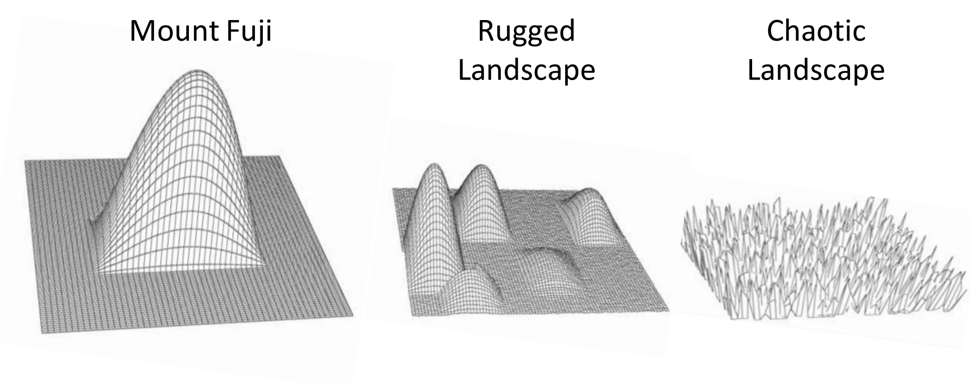

Nature of the problem:	Simple	Complicated	Impossible
Search Algorithm:	Hill-climbing	Evolution: variation, selection, replication	Random
Quality of the solution:	Optimal	Local peaks Sub-optimal	Low

Figures produced using QtiPlot.

Figure 2 – Epistatic Couplings and Internal Constraints

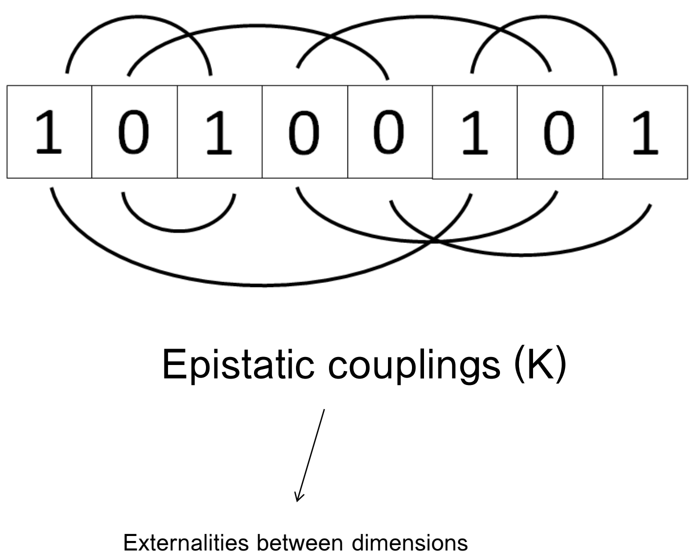

Figure 3 – Beliefs and the actual relationship between institutions and outcomes

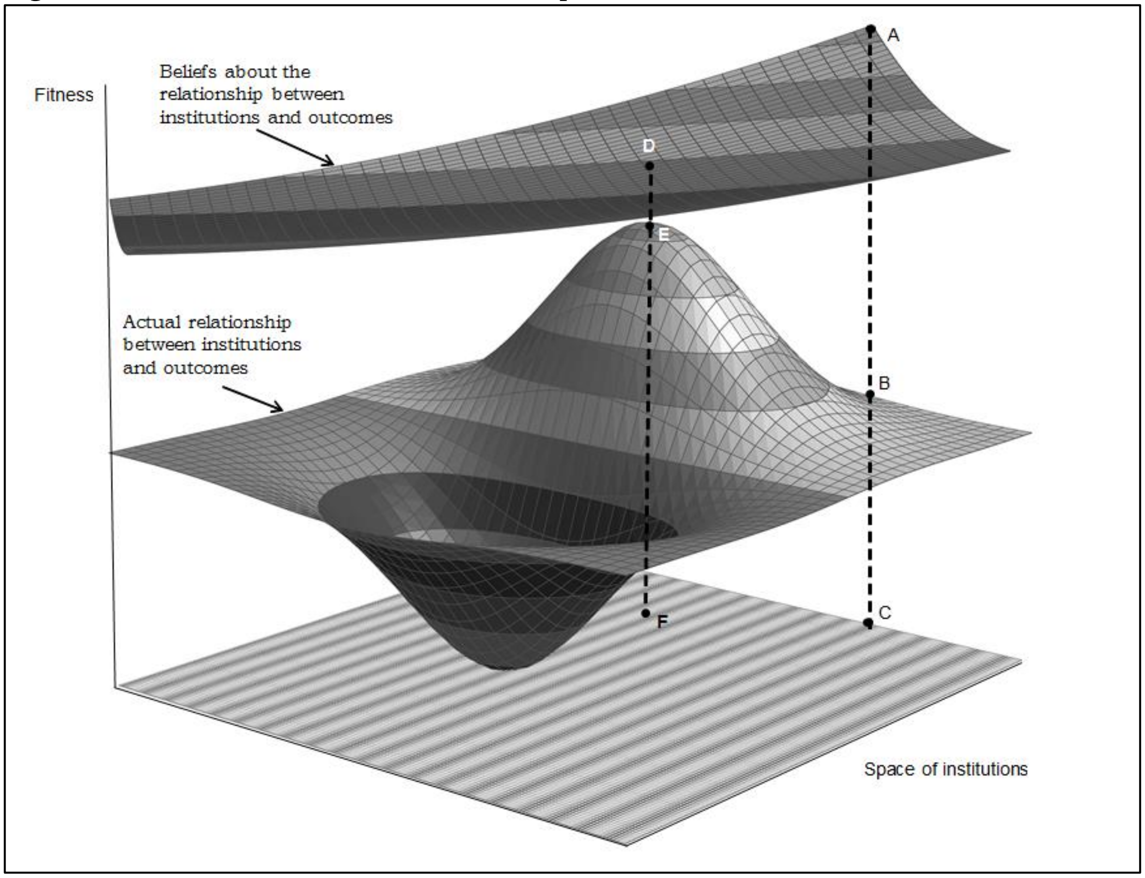

Figure produced by the author.

Figure 4 – Coevolution of beliefs and technology

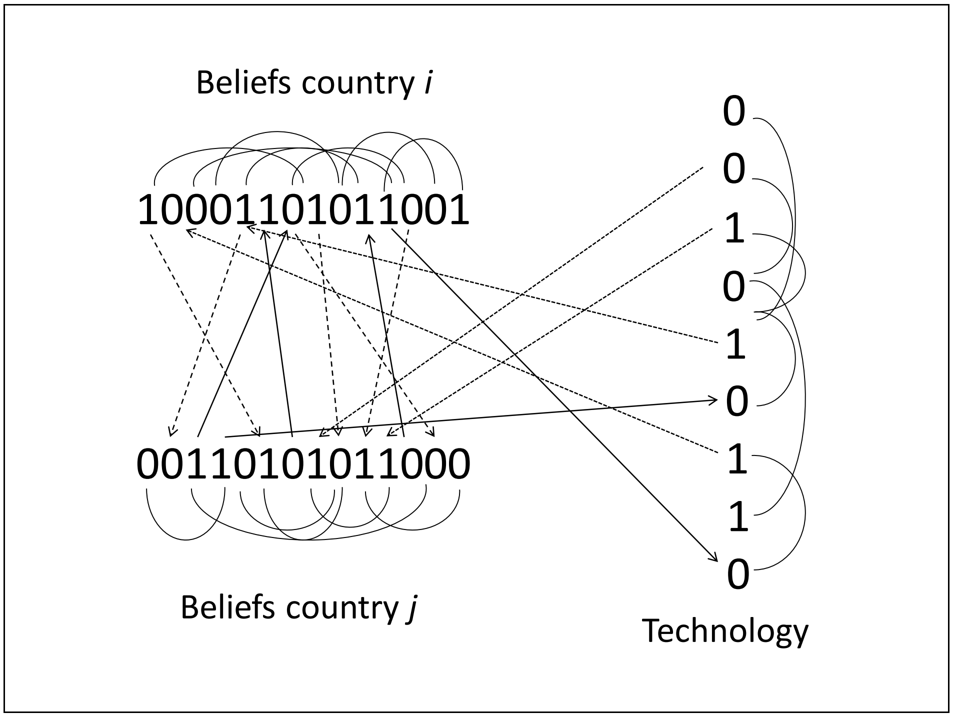

Figure 5 – Hypothetical Institution Space and Fitness Function for Development

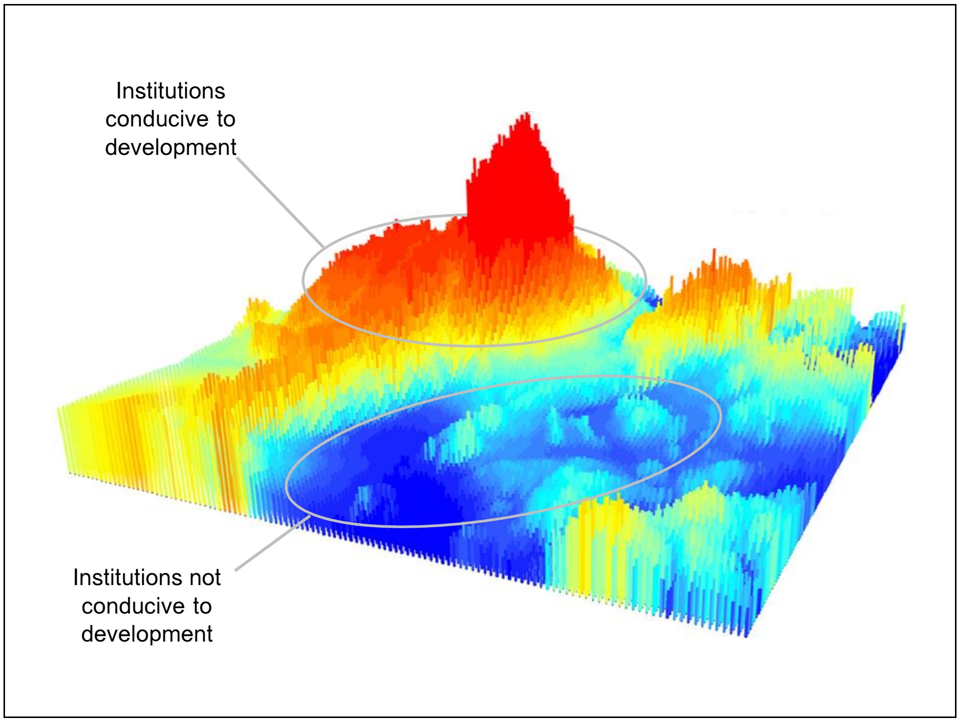

Fitness landscape figure from Weaver and Knight (2014: 511).