Complex Adaptive Systems and the Threshold Effect: Views from the Natural and Social Sciences:
Papers from the AAAI Fall Symposium
(FS-09-03)

Formal Measures of Dynamical Properties: Tipping Points

Aaron L Bramson
Center for the Study of Complex Systems
University of Michigan
321A West Hall
1085 S. University Ave.
Ann Arbor, MI 48109-1107

Abstract

To help realize the potential of complex systems models we need new measures appropriate for capturing processes that exhibit feedback, nonlinearity, heterogeneity, and emergence. As part of a larger research project encompassing several categories of dynamical properties this paper provides formal and general definitions of tipping point-related phenomena. For each tipping concept this paper provides a probabilistic definition derived from a Markov model representation. We start with the basic features of Markov models and definitions of the foundational concepts of system dynamics. Then several tipping point-related concepts are described, defined, measured, and illustrated with a simplified graphical example. The paper finishes with several branches of future work involving new measures for complex systems and the fusion of research domains.

1. Introduction

The idea of tipping points has captured the public's attention from topics as diverse as segregation, marketing, political unrest, material science, ecosystem stability, and climate change. However the concept of a tip has not been generally and formally defined and, as a result, the term's uses across these various applications are inconsistent. At times a tipping point refers to a threshold beyond which a system's outcome is known. Other times a tipping point describes an event that suffices to achieve a particular outcome, or an aspect of such an event, or the time of such an event. Another use of the term 'tipping point' is to label the conditions to which the system is most sensitive. The idea is frequently tied up with processes such as positive feedback, externalities, sustainable operation, perturbation, etc. but not in ways that explicitly draw the connections. This paper aims to elucidate the distinctions among these and other uses of the term 'tipping point' based on features of system behavior that are independent of the substantive domain.

To accomplish this conceptual analysis the paper puts forth a formal definition for each concept to measure the associated property. The analysis utilizes Markov model representations of systems constructed in a particular way. The details of creating the necessary Markov representation are explained in detail in other work (Bramson 2009) as are the

formal definitions for familiar landmarks in system dynamics. These features will be only briefly covered here. With this foundation in place various tipping point-related concepts are described, defined, and illustrated with simplified graphical examples.

2. Background

Much previous work in finding and measuring properties of system dynamics has focused on explanation - on answering the *why* questions. Not surprisingly since these discussions were couched in scientific contexts where a particular phenomenon required explanation. Each such jaunt into explaining tipping points was accompanied by a custom-suited methodology capable of generating and detecting that property within the model provided (to answer the *how* question). Each such measure was designed to explain the dynamical property within a particular model and domain. This approach is methodologically limiting because in order to explain how a process generates tipping behavior one has to model that process explicitly and develop a new measure for each model.

The current work is one of pure methodology: It is meant to be completely abstract and general and therefore capable of measuring these system properties in any system. This paper provides techniques to answer the *whether* and *how much* questions which are key to informing the *how* and *why* questions. A general methodology provides a framework through which all modelers (and some data analysts) can determine whether and how much tippiness obtains . . . and compare results across models regardless of the generating mechanisms. The ability to compare measures across systems is achieved through a focus on scale-free measures - measures that do not depend on the size of the system being analyzed.

2.1 Tipping Points

The term 'tipping point' was first coined by Morton Grodzins in 1957 (Grodzins 1957) to describe the threshold level of non-white occupants that a white neighborhood could have before "white flight" occurred. The term continued to be used in this context through the work of Eleanor Wolf (Wolf 1963) and Thomas Schelling (Schelling 1971) who also extended the concept to other similar social phenomena. Though these researchers had a specific usage with

narrow focus, the idea of a critical parameter value past which aggregate behavior is recognizably different spread across disciplines where its meaning and application varied considerably.

Malcolm Gladwell's pop sociology book *The Tipping Point* (Gladwell 2000) has played a significant part in bringing the term to the public's awareness. The notion of tipping point most frequently used by Gladwell is an event that makes something unusual (such as Hush Puppy shoes) become popular. More precisely this is a critical value for producing a phase transition for percolation in certain heterogeneous social network structures. This form of tipping point behavior also appears in the work of Mark Granovetter (Granovetter 1978) and Peyton Young (Young 2003) for the propagation of rioting behavior and technology respectively. This version of tipping will play only a minor role in what follows, however the fact that the expression has made it into the everyman's conceptual vocabulary boosts the importance of establishing rigorous scientific definitions and usage.

Tipping points have also appeared as a recent trend in reports of climate change. James E. Hansen has claimed that "Earth is approaching a tipping point that can be tilted, and only slightly at best, in its favor if global warming can be limited to less than one degree Celsius." (Farrell 2006) This usage reflects Hansen's belief that "Humans now control the global climate, for better or worse." Gabrielle Walker states, "A tipping point usually means the moment at which internal dynamics start to propel a change previously driven by external forces." (Walker 2006) It is unclear whether Walker's and Hansen's comments are compatible; and even if their usages are meaningful within their fields they fail as general characteristics. Identifying tipping points (as a property of system dynamics) should not depend on whether humans are in control of system behavior or what is driving the dynamics.

But not all heretofore definitions of the term 'tipping point' have been loose or subject-matter specific. It is often deployed as a semi-technical term in equation-based models of various sorts. For example, it can refer to an unstable manifold in a differential equation model, the set of boundary parameters for comparative statistics (Rubineau 2007), or inflection points in the behavior of functional models. But not all models can be faithfully represented as systems of equations and this limits the usefulness of equation-dependent definitions.

2.2 Markov Modeling, Network Theory, Graph Theory, etc.

Markov modeling has a long history in mathematics, engineering, and in applications to fields as diverse as condensed matter physics, genomics, sociology, and marketing. Techniques to use Markovian processes to uncover information about system dynamics fall within the field of ergodic theory. In their abstract form one can compute features such as the equilibrium distribution, expected number of steps between two states (with standard deviation), reversibility, and periodicity. These features gain added meaning when interpreted for the system being modeled, but this paper utilizes them as part of defining general properties of system dynamics.

Computer scientists have long been analyzing networks in the form of actual communication networks as well as various abstractions from these problems. They have invented several useful measures and exceptionally well-crafted algorithms to calculate connectedness, load-bearing properties, path switching, transmission speed, and packet splitting and fusion to name a few. My analysis borrows heavily from this work in terms of algorithms, though each has been repurposed to the abstract Markov model system representation. Hardware engineers and their physicist partners have worked out several interesting measures for circuit design problems. Multiple paths, variable resistance, flow injection, capacitance and many other characteristics of electronic circuits have analogs in the Markov models presented below.

And finally graph theory offers a few useful measures for our purposes, and moreover provides a wealth of definitions for graph structure and node relationships. Structural properties will play a larger role in future work addressing changes of resolution and in establishing equivalence classes of system dynamics (see future work).

3. Motivations and Applications

The measures defined here are meant to stand on their own as improvements in our conceptual understanding of the included features of system dynamics. By differentiating and formally defining these properties of processes we gain both a common vocabulary with which to discuss our models and a detailed typography of behavior to include and detect in our models. Many of the applications I have in mind are to include these measures in constructive models across multiple disciplines where the models are iterated with multiple initial settings and/or have stochastic parameters. These include game theoretic models, network models, physical models, and the whole gamut of models which may be considered agent-based. The application to agent-based models of complex adaptive systems is the most important since this is where existing measuring techniques fall the shortest of fulfilling basic needs. One major goal of complex systems research is to identify common underlying mathematical properties in a myriad of seemingly very different phenomena. The Markov modeling technique used here allows us to create a common representation of almost any system's dynamics and the measures developed here and elsewhere thus provide immediate ways to compare the dynamical properties of systems across domains. In addition to generative models and simulations certain static data sets - the sorts collected by surveys - are also analyzable via this methodology if the data satisfy certain criteria. Readers are referred to (Bramson 2009) for more information about creating the Markov model from collected data.

4. Defining a Markov Model

In consideration of limited space and because a detailed account of how to create the Markov model representation is provided elsewhere (Bramson 2009) I will only present a brief overview of this mathematical structure to highlight

those features which will be key to the tipping point analysis at the heart of the current paper. Fundamentally a Markov model is simply a collection of states and a set of transition probabilities between pairs of states. Different applications of Markov models take different system features as the states, but the nodes in the Markov models used here represent a complete description of a state of the system (see below). The transition probabilities represent either observed system dynamics or theoretically posited state changes. Given that states and transitions are defined this way the set of states and their transition probabilities are constant for the Markov models utilized in this paper.

A system state is a complete set of instantiations of the aspects of the system (values for variables, existence for agents, etc.). Throughout we will analyze systems with a finite (but possibly arbitrarily large) number of states each with a finite number of aspects.

Definition 41. A state in the Markov model is a complete specification of the aspects of one configuration of the system.

Definition 42. Two states are represented as one state of the Markov model if all the aspects of the two states are identically valued.

Example 41. If our system is an iterated strategic-form game played by six players
$$P_i \in \{P1, P2, \ldots, P6\}$$
each with four possible actions
$$a(P_i) \in \{a1, a2, a3, a4\}$$
then each state of the system has six aspects and each aspect takes on one of four values. That is
$$S_i = \{a(P1_{(i)}), a(P2_{(i)}), \ldots, a(P6_{(i)})\}$$
and a particular state S_3 might be $\{a3, a2, a3, a1, a4, a3\}$. There are $6^4 = 1296$ combinations of four actions for six players, but the Markov model may not include all of them. Recall that the model is expected to be built from either collected data or a theoretical model so some combinations of aspect values may be unobserved, theoretically impossible, or simply irrelevant.

A set of n states is demarcated with boldface type: $\mathbf{S} = \{S_1, S_2, \ldots S_n\}$. The set of all the states in the Markov model is $\mathbf{N}$ which has size $|\mathbf{N}| = N$; thus N is also the number of nodes in the graph representation. The state of the system at time t (denoted s_t) changes to s_{t+1} in discrete, homogenous time intervals. State transitions are probabilistic and specified by the system's transition diagram or matrix. We write the probability of transitioning from state S_i to state S_j as $P_{ij} := P(s_{t+1} = S_j | s_t = S_i)$. It will later be useful to denote the set of transitions $\mathbf{E}$ and the size of this set as $|\mathbf{E}|$.

5. Special States and Sets

To use a Markov diagram to represent system dynamics we will need to define various types of system behaviors in terms of system states, sets of states, and state transitions. As a preliminary to the common features of system behavior I will present definitions of some structural features that will be utilized.

5.1 Paths

In graph theory a path (of length ℓ) is typically defined as a set of vertices and edges satisfying the schema $v_0, e_1, v_1, e_2, \ldots, v_{\ell-1}, e_\ell, v_\ell$ where the edge e_i links the vertex v_{i-1} to v_i (Lint 1992). Self-transitions, which represent both a lack of change and a change too small to count as a state change, are an important feature of Markov modeling and hence both nodes and edges may be repeated along paths. So 'path' as it is used here is the broader notion sometimes called a 'walk' in the graph theory literature. A path in a Markov model could be defined as a set satisfying the same schema used in graph theory, but we will use a slightly different definition to make the probabilistic aspects explicit.

Definition 51. A *path* is an ordered collection of states and transitions such that from each state there exists a positive probability to transition to the successor state within the collection. A path from S_i to S_j is denoted $\tilde{\mathbf{S}}(S_i, S_j)$ or $\widetilde{S_i \, S_j}$

To specify intermediate states (*markers*) for the system to pass through we can write $\widetilde{S_i \, S_j \, S_k}$ to denote a path from S_i to S_k that passes through (at least) S_j. To specify a long sequence of path markers this paper uses the notation $\tilde{\mathbf{S}}(S_0, S_1, \ldots, S_T)$.

To completely specify each state along a path we adopt the notation $\overrightarrow{S_i \, \ldots \, S_j}$ for short sequences and $\vec{\mathbf{S}}(S_0, S_1, \ldots, S_T)$ for long ones.

Definition 52. The *length* of a path is the number of transitions taken between the first and last states.

$$\ell(\vec{\mathbf{S}}(S_i, \ldots, S_j)) := \sum_{t=0}^{T-1} |\{\overrightarrow{s_t \, s_{t+1}}\}|$$

This (possibly overly complicated) formal definition of length simply uses features of the definition of path above, but it is equivalent to the number of edges traversed along the path.

Definition 53. A *cycle* is a path that starts and ends with the same state.

$$\widetilde{S_i \, S_i}$$

Graph and network theorists have developed a great many algorithms for finding paths, calculating their lengths, and measuring properties germane to their application in those fields. Some of those will come up later in measuring properties of system dynamics, but the definitions and theorems presented above will suffice to move forward in examining our Markov models.

5.2 Landmarks in System Dynamics

This subsection provides definitions for common structural properties of Markov models used to represent system dynamics. Many of these features have existing definitions in terms of matrix operations or limiting distributions, but this paper will present probabilistic definitions based on the finite state transition representation. My motivation for the alternative definitions is to facilitate clear intuitions about a system's processes and how to measure them precisely.

Definition 54. A system state that always transitions to itself is called an *equilibrium* or *stable state*.[1] An equilibrium e_i is a state S_i such that

$$P(s_{t+1} = S_i | s_t = S_i) = 1.$$

In some cases a set of states plays a role similar to that of an equilibrium.

Definition 55. An *orbit* is a set of states such that if the system enters that set it will always revisit every member of the set and the system can never leave that set. $\mathbf{S}$ is an orbit if

$$\forall i \; \nexists h \; P(s_{t+h} = S_i \in \mathbf{S} | s_t \in \mathbf{S}) = 0$$

and

$$\forall i \; \forall h > 0 \; P(s_{t+h} = S_i \notin \mathbf{S} | s_t \in \mathbf{S}) = 0.$$

Definition 56. An *oscillator* is an orbit that is also a cycle.

Definition 57. An *Attractor* (denoted A_i) is either an equilibrium state or an orbit of the system.

The choice of a resolution determines whether an orbit appears as an equilibrium or *vice versa*. In cases where attractor avoidance is the aim of the model (as is often the case for complex systems) we can collapse orbits into a single attractor state without loss of information. For this reason I will use 'A_i' as if it were a single state except in cases where it being an orbit affects the analysis.

For many systems an equilibria analysis is inappropriate. It is not the case that these systems fail to have attractors, it's just that the goal of such systems is to remain in continual flux and avoid equilibria and other "point attractors" (i.e. attractors that incorporate a small percentage of the total number of states). Given the definitions established above every Markov model must have some set of states satisfying the conditions for being an attractor, but this may be the whole set of states seen in the dynamics. In many cases the important concept is that of a core (see below) which can be thought of as quasi-attractors.

Definition 58. Those states from which the system will eventually move into a specific attractor are said to be in that attractor's *basin of attraction*. The basin of A_i or $\mathbf{B}(A_i)$ is a set of states $\mathbf{S}$ such that

$$\exists h \geq 0 \; P(s_{t+h} = A_i | s_t \in \mathbf{S}) = 1$$

Some systems may spend a great deal of time in a basin of attraction before reaching the attractor located within it thus making system behavior in the basin similar to an orbit itself. In such cases it is sometimes helpful to utilize the following property to describe and make inferences about system behavior.

Observation 51. *Once in a basin of attraction the system can never leave it.*

$$\forall h \; P(s_{t+h} \in \mathbf{B} | s_t \in \mathbf{B}) = 1$$

---

[1]As is defined formally in the paper on robustness measures (Bramson 2009), stability refers to a tendency to self-transition. Hence an equilibrium is equivalent to a **fully** stable state.

Definition 59. The *support* of a state (also known as its *in-component*) is the set of states which have a path to it. The support of S_i or $\mathbb{S}(S_i)$ is the set of states such that

$$\forall j \; S_j \in \mathbb{S}(S_i) \Rightarrow (\widetilde{S_j \; S_i}).$$

We can expand this definition to the support of a set of states $\mathbf{S}$ as the union of the supports of the members of $\mathbf{S}$.
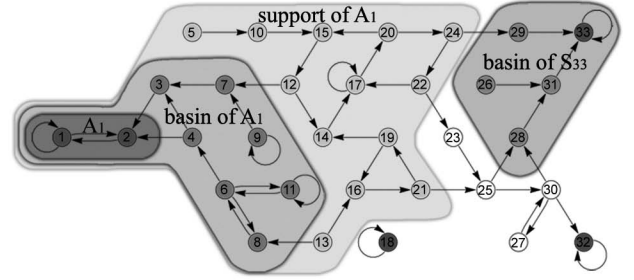

Figure 1: Attractors, Basins of Attraction, and Support

The *indeterminate states* of a system, ones that are not members of a basin of attraction, convey a wealth information about the system's dynamics and its future states. Recall that the whole system may be a single orbit and there may be no indeterminate states. But if there are multiple attractors, then the indeterminate states are the ones in multiple attractors' supports.

Definition 510. The *overlap* of a collection of states (whether attractors or not) is the set of states in all of their in-components (i.e. the intersection of supports). The overlap of $\{S_i, \ldots, S_j\}$, written $\mathbf{\Omega}(S_i, \ldots, S_j)$, is the set of states in

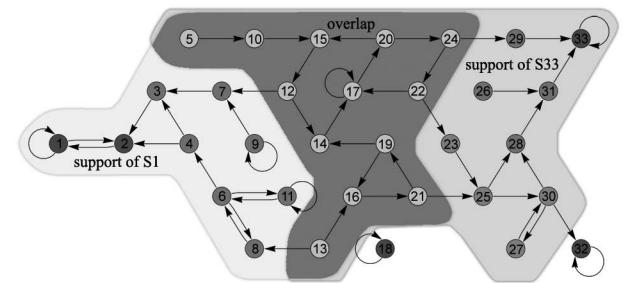$$\bigcap_{S_g \in \{S_i, \ldots, S_j\}} \mathbb{S}(S_g)$$

Figure 2: Supports and Their Overlap

The overlap states are of interest because these are the states with positive probabilities for ending up in each of the attractors for which the supports are overlapping.

Observation 52. *If no attractors' supports overlap then every attractor's support is just its basin of attraction and the system has a deterministic outcome.*

Definition 511. A state's *out-degree* is the number of distinct successor states (states that may be immediate transitioned into). The out-degree k_i of state S_i equals

$$|\{S_j : P(s_{t+1} = S_j | s_t = S_i) > 0\}|$$

S_k will be used to denote a neighboring state and $\mathbf{S_k}$ the set of neighboring states.

Definition 512. The *reach* of a state (also called its *out-component*) is the set of states that the system may enter by following some sequence of transitions; i.e. *all possible* future states given an initial state. The reach of S_i or $\mathbf{R}(S_i)$ is the set of S_js such that

$$\exists h > 0 \ P(s_{t+h} = S_j | s_t = S_i) > 0$$

Theorem 53. *Every successor state's reach is less than or equal to the intial state's reach.*

$$\forall i, j \ \overrightarrow{S_i \, S_j} \Rightarrow \|\mathbf{R}(S_i)\| \geq \|\mathbf{R}(S_j)\|$$

Theorem 53 generalizes to all paths (which is just a sequence of transitions) so that reach never increases as the systems transitions along any path. This property relies on the fact that the transition structure is fixed for the Markov models used in this methodology.

A *strongly connected component* of a directed network is a set of vertices such that there is a path from every vertex in the set to every vertex in the set (including itself). We will find the same concept useful, but this paper adopts a different name for it.

Definition 513. A *core* of a set is a subset wherein every member of the subset is in the reach of every member of the subset. The core of some set $\mathbf{S}$ is written $\mathbb{C}_\mathbf{S}$ and is a subset satisfying the condition

$$\bigcap_{S_i \in \mathbf{S}} \bigcap_{S_j \in \mathbf{S}} S_j \in \mathbf{R}(S_i)$$

Some sets will have multiple cores – the set of $\mathbf{S}$'s cores can be called $\mathbf{S}$'s *mantle*.

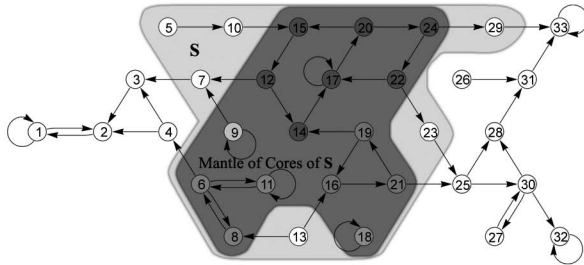**Observation 54.** *Every state in a core has the same reach.*

Figure 3: The Cores of an Arbitrary Set

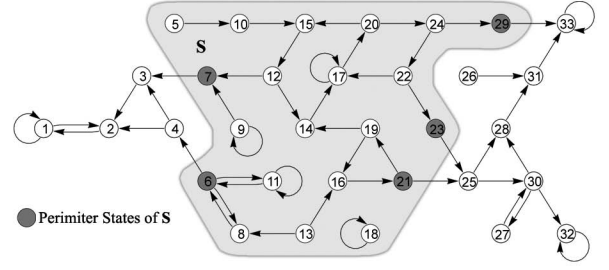The states on the boundary of a set are a useful set of states to identify.

Figure 4: The Perimeter States of an Arbitrary Set

Definition 514. The *perimeter* of a set, $\mathbb{P}(\mathbf{S})$, is a collection of those states in the set that may transition to states outside the set. That is, $\mathbb{P}(\mathbf{S})$ is the set of states such that

$$P(s_{t+1} \notin \mathbf{S} | s_t \in \mathbf{S}) > 0$$

In keeping with the core and mantle analogy, the perimeter states of the mantle of $\mathbf{S}$ will be referred to as $\mathbf{S}$'s *crust*. Perimeter states themselves, without further specification, describe one commonly deployed (though weak) concept of tipping points, although we will see in the next section that specifying different base sets produces different types of tips.

In the next section we apply the above definitions in different combinations and different contexts to identify various system behaviors. Most extant systems analysis focuses on equilibria, but a lot of interesting behavior happens away from equilibrium. In the indeterminate states of a system we cannot know precisely which states the system will reach or which state it will be in at a given time in the future, we only know probability distributions over the future states. But by understanding a system's behavior we might know whether some particular change facilitates a specific outcome or path through the dynamics - this may be helpful information. Discerning these sorts of facts about system dynamics can increase one's information about the system (in both the technical and colloquial senses). Considerations such as these are the building blocks of the formal theory of tipping point-related phenomena immediately to follow.

6. Tipping Phenomena and Related Concepts

Using Markov models and the states and sets defined above as a springboard, this section defines and briefly describes several terms related to the concept (or more to the point, concepts) of tipping points. Critical phenomena and tipping points of various kinds share the defining feature that (for whatever reason) behavior is different before and after some transition. Behavior in this analysis is just the properties of system dynamics. There are, of course, many ways in which the properties of system dynamics can differ and each way is a different kind of tipping phenomenon.

6.1 Levers

Recalling that state changes occur if and only if there is a change in some aspect of the initial state, our analysis of tipping phenomena starts with state aspects.

Definition 61. The *levers* of a state are the aspects of a state such that a change in those aspects is sufficient to change the system's state. The levers of S_i, denoted $L(S_i)$, is

$$\bigcup_{j=1}^{k} \bigcup_{h=1}^{Q} X_{h(i)} \neq X_{h(j)}.$$

So given a state in the Markov model, the levers are those aspects of the state that are different in any neighboring state and each such aspect X_h is a distinct lever. In some cases neighboring states will differ by more than one aspect. In those cases the respective element of the set $L(S_i)$ will be a list of all the aspects that need to change as one lever.

Definition 62. A *lever point* is a transition resulting from a change in a particular aspect (or set of aspects). An aspect's lever points is the collection of transitions that a change in that aspect (or those aspects) alone generates. The lever points of X_h is the set of transitions created by

$$\bigcup_{\overrightarrow{S_i \, S_j}}^{\mathbf{E}} X_{h(i)} \neq X_{h(j)}$$

Levers and lever points work complementarily: for levers we pick a state and find the aspects that change and for lever points we pick the aspect change and find the state changes it produces.

It is occasionally helpful to refer to the aspect change(s) that generate a specific transition.

Definition 63. $L(\overrightarrow{S_i \, S_j})$ symbolizes the *lever set* of $\overrightarrow{S_i \, S_j}$: the aspect or aspects that differ between S_i and S_j.

In some applications we will be interested in how many aspects change for a transition.

Definition 64. The *magnitude* of the lever set of a specific transitions is $|L(\overrightarrow{S_i \, S_j})|$.

Though levers as they are defined here do not depend on the ability to control that aspect, the choice of 'lever' for this concept is motivated by the realization that in some models control of some aspects is available. One may be performing a tipping point analysis precisely because one is choosing levers to bring about one state versus another (or agents within the model may be choosing).

Example 61. Imagine a model wherein each aspect is a variable representing some part of a policy (e.g. amount of money spent on each line item). Each aspect change has an associated cost (legal, bureaucratic, time, etc.). The modeler may be trying to determine the lowest cost, feasible route from the current policy to some desired policy; or perhaps to determine how far policy can be changed on a specific budget. The cumulative magnitudes of lever sets along a path may adequately approximate such a cost measure. In general the sum of the magnitudes along a path is a rough measure of how difficult it is for the system to behave that way. Techniques from circuit design applied to the Markov model may be gainfully applied to such models.

In some contexts we may wish to know how much change an aspect is responsible for across the system's dynamics.

Definition 65. The *strength* of a lever is the sum of the probabilities of all transitions that result from changing that lever. The strength of X_h is equal to

$$\sum_{\overrightarrow{S_i \, S_j}}^{\mathbf{E}} P(\overrightarrow{S_i \, S_j} | X_{h(i)} \neq X_{h(j)}).$$

This measure is not scale-free since the sum depends on the number of transitions in the Markov model, but it is useful for comparing levers within a system. The strength measure could be used, for example, to determine which aspects to control to maximize (or minimize) one's ability to manipulate the system. It could also be associated with a cost of letting that aspect vary over time. We will revisit levers below in other forms as they apply to other measures of system dynamics.

6.2 Thresholds

In some cases we are interested not just in which aspects change through a transition but also in the **values**[2] of levers at transitions.

Definition 66. A *threshold* or *threshold point* is a particular value for a lever such that if the value of the aspect crosses the threshold value it generates a transition. So x is a threshold value of $\overrightarrow{S_i \, S_j}$ if

$$X_{h(i)} = x \text{ and } X_{h(j)} \neq x.$$

This definition can be applied *mutatis mutandis* for a set of values $\{x\}$ for a lever set which can be distinguished by the name *threshold line* when appropriate.

If there are multiple states with transitions crossing the same threshold value then knowing that information refines our understanding of the lever's role in system dynamics. Thus determining the threshold value for one transition is merely a means to the end of determining the strength of the levers with that threshold.

Definition 67. The *threshold strength* of x is the strength of the levers for which x is the threshold value:

$$\sum_{\overrightarrow{S_i \, S_j}}^{\mathbf{E}} P(\overrightarrow{S_i \, S_j} | X_{h(i)} = x \text{ and } X_{h(j)} \neq x.$$

If a particular value for a particular aspect plays a large role in system dynamics then crossing that threshold is another oft-used version of tipping point behavior.

These definitions of threshold and threshold strength only require that the end state's value be different from the start state's value. In common usage, however, thresholds establish different and separate boundary values for ascending and descending values. If a threshold only affects system dynamics in one direction then we can determine that from the Markov model using the following definitions.

---

[2]Recall that many of the things that can be included as aspects of states are not numeric parameters and so what counts as a "value" for that aspect is meant to be interpreted broadly.

Definition 68. An *upper bound threshold* of $\overrightarrow{S_i\,S_j}$ is a value x such that

$$X_{h(i)} = x \text{ and } X_{h(j)} > x.$$

Definition 69. A *lower bound threshold* of $\overrightarrow{S_i\,S_j}$ is a value x such that

$$X_{h(i)} = x \text{ and } X_{h(j)} < x.$$

The threshold strength measure can be adapted to these ascending and descending definitions in the obvious ways. Sets satisfying these definitions can tell us how frequently crossing that threshold in that direction acts as a lever. If there are multiple states with transitions crossing the same threshold value then knowing that information refines our understanding of the lever's role in system dynamics.

Example 62. This general definition admits examples from many different kinds of systems and can even apply to parts of systems (such as agents). In Granovetter's model of riot spreading (Granovetter78) we can talk of each agent having its own threshold - the number of rioting agents necessary to make each agent join the riot. This is just the same threshold definition applied to a lever set where the levers happen to be the same feature of each agent. In Granovetter's model the threshold value is the same in both directions.

We can also talk of thresholds in the properties of the system dynamics that track how the system transitions through states. Instead of being a value for an aspect within the model, it would be a value for one of the measures defined for properties of system dynamics.

6.3 Critical Behavior

In some system analyses the property of interest is what is available for the future . . . in the most general terms. If one does not know much about a system's dynamics then even knowing how many states could potentially be transitioned to (i.e. the size of the reach) provides an informational benefit. The measures below become increasingly refined and detailed, but we start with some simple measures that may suffice for some applications (and may be the best available).

Definition 610. A state's stretch is the number of states in its reach. So the stretch of S_i equals

$$|\mathbf{R}(S_i)|.$$

Definition 611. A system dynamic (i.e. a particular state transition) is considered *critical behavior* if and only if it produces a decrease in stretch; that is, critical behavior is any $\overrightarrow{S_i\,S_j}$ such that

$$|\mathbf{R}(S_i)| > |\mathbf{R}(S_j)|.$$

In addition to identifying the transitions that limit the system's future states, we can also measure how critical the transition is. Subtracting the end state's stretch from the start state's stretch provides such a measure, but it is not scale-free (because both the range of values and the particular value for this measure depends on the total size of the system) and so cannot be readily compared across different systems. We can normalize the stretch difference with the size of the system to which it is being applied to produce a percentage measure.

Definition 612. The *stretch-gap* of a transition is the change in the percent of the total number of states that can be reached. This quantity equals

$$\frac{|\mathbf{R}(S_i)|}{N} - \frac{|\mathbf{R}(S_j)|}{N}.$$

Because this measure includes the total number of states in the system it clearly is not scale-free either. However despite this limitation it does provide information about the system's future and is an intuitive way to compare transitions within the same system - even at different resolutions. As example 63 below demonstrates the stretch-gap reports how much of the system's state space is cut off by each transition and this information could be used, for example, to manipulate system dynamics to prolong system longevity. We also have an alternative, fully scale-free, measure of the drop in reach across a transition.

Definition 613. A **transition's** *criticality* is one minus the ratio of the start and end states' stretch. The criticality of $\overrightarrow{S_i\,S_j}$ equals

$$1 - \frac{|\mathbf{R}(S_j)|}{|\mathbf{R}(S_i)|}.$$

Recall from theorem 53 that from any initial starting point, as the system transitions through its states the sizes of the states' reach are monotonically decreasing. As a result of that theorem we have the following corollary regarding the range of values for the ratio of reaches.

Corollary 61. *A transition's criticality will be between zero and one.*

Proof. Let $a = |\mathbf{R}(S_j)|$ and $b = |\mathbf{R}(S_i)|$ for transition $\overrightarrow{S_i\,S_j}$. From theorem 53 $a \leq b$. $a = b$ produces $\frac{a}{b} = 1$ which yields a criticality of zero. For $a < b$ we can decrease a or increase b to find the other bound, but since a and b are natural numbers increasing b is the better approach. Using a well-known mathematical fact suffices for finding the other bound: $\forall a \ \lim_{b \to \infty} \frac{a}{b} = 0$. $\qquad\square$

Observation 62. *Transitions within a cycle (which includes self-transitions) always have zero criticality and zero stretch-gap.*

The concept of criticality agrees with this measure insofar as any transition that has no affect on what states may be visited in the future should not be a critical transition. Critical transitions, or the states preceding them, are another oft-used tipping concept.

Example 63. The system represented as figure 5 has thirty-three states in total. Each one is coded by its stretch. Compare the patterns in shading to the attractors, basins, and support in figure 1 and the overlap in figure 2. Stretch alone, though a rather simplistic measure, works decently to partition the system dynamics into regions of similar behavior. Stretch-gap performs well as a discriminator that groups states into these regions in a way similar to spectral analysis for detecting community structure in networks (Clauset 2008). To wit, stretch drops between zero and four states within any basin or overlap, but drops four to eighteen states
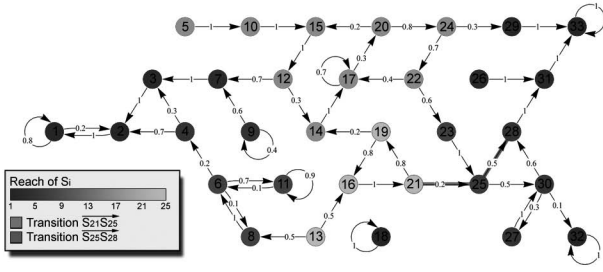
8

Figure 5: Stretch-Gap and Criticality Measures

crossing a boundary (in this example). This is not completely reliable, of course, but for many systems this simple technique may provide all the information required; and it may be the best one can do with available data.

S_{21} has a stretch of 21, S_{25} has a stretch of 6, and S_{28} has a stretch of 2. There are 33 states in this system so the stretch-gaps of $\overrightarrow{S_{21}\,S_{25}}$ and $\overrightarrow{S_{25}\,S_{28}}$ are 45.45% and 12.12% respectively. That means that 45.45% fewer of the system's states can be reached after the $\overrightarrow{S_{25}\,S_{28}}$ transition. We can also use this to determine the stretch-gap of $\overparen{S_{21}\,S_{28}}$ as 57.58% regardless of the particular path taken. Only the start and end states' stretches are necessary to calculate this, but the result is always equal to the sum of the stretch gaps of each transition taken.

Let's compare these figures to the criticality of the same transitions. The criticality of $\overrightarrow{S_{21}\,S_{25}}$ is 0.714 and the criticality of $\overrightarrow{S_{25}\,S_{28}}$ is 0.667. That means that the system only has 71.5% of the possible future states in state S_{25} as it did in S_{21}. A composite measure is also possible for the criticality of $\overparen{S_{21}\,S_{28}}$. It can be calculated just using the start and end states' stretches, using the standard percentage of a percentage of a percentage . . . calculation.[3] So the criticality of $\overparen{S_{21}\,S_{28}} = 0.714 + 0.667 \cdot (1 - 0.714) = 0.905$.

These measures above are intended to be just rough measures useful in certain limited contexts and when information about the system is limited. For starters, these measures consider only the structure of the Markov models, not the probabilities. Also, they apply to transitions rather than states. Both limitations are now overcome with richer definitions of the relevant dynamical properties.

Definition 614. The *criticality* of a **state** is the probabilistically weighted sum of the criticality of all the transitions from that state. So to find the criticality of S_i we calculate

$$\sum_{j=1}^{k} P_{ij}\left(1 - \frac{|\mathbf{R}(S_j)|}{|\mathbf{R}(S_i)|}\right)$$

where by convention of the use of k that is the sum over S_i's neighbors.[4]

---

[3]The value equals the iterated sum of the previous transition's criticality and the product of the transition criticality with the previous transition's criticality's complement

[4]In this case it does not matter whether the sum is limited to

Because by definition of a Markov model the sum of the exit probabilities sum to one, state criticality will also be a scale-free measure with values between zero and one. All these criticality measures quantify the constriction of future possibilities on a state-by-state basis which is useful if we want "to keep our options open". Sometimes that is exactly what we want to measure, but sometimes we will want to measure system dynamics with reference to some particular features and that is what the following definitions for tipping points allow us to measure.

Critical Levers As a refinement of levers from definition 61 we can apply the lever concept to critical states to identify another feature of system dynamics.

Definition 615. A state has a *critical lever* if a change in that aspect (or those aspects) of the state will reduce the reach.

This merely combines the concept of a lever with the concept of critical behavior (definition 611). By looking more deeply at the aspects driving the state changes and calculating the magnitude and strength of different critical levers we can gain a better understanding of how microfeatures generate the macrobehavior of the model.

6.4 Tipping Behavior

As mentioned in the section introduction, the common feature of the measures in this section is that some states or transitions mark a shift in the properties of a system or of its states. For the criticality measures above the difference was the number of reachable states. The following measures generalize to any sets distinguished by a chosen characteristic. Given states exhaustively compartmentalized (i.e. partitioned) by the property (or properties) of interest the following techniques can find where shifts occur and measure their magnitude.

Tipping Points For some models we are interested in the achievement of a particular state (e.g. an equilibrium) or a particular system behavior (e.g. a path linking two states). We denote the particular state (or set) of interest as the *reference state* (or *reference set*). Below we will see examples of specific reference states (e.g. attractors and functional states) but first the general case. There are many ways in which behavior may change with respect to a reference state or set (e.g. probability of reaching it, probability of returning to it, or probability of visiting an intermediate state): each property may partition the states into different equivalence classes (groups with the same value of the property). It is the movement between equivalence classes that counts as tipping behavior.

Definition 616. A *tipping point* is a state which is in the perimeter of an equivalence class for some property.

Recall from definition 514 that perimeter states are those from which the system's dynamics can leave the specified set. Because the sets here are determined by the properties of system behavior, leaving a set implies a change in that

---

span over neighbors or all the vertices because $P_{ij} = 0$ for S_j that are not neighbors. This convention will be used throughout - including cases where limiting an operation to neighbors matters.

behavioral property - and that is a tip. This definition does not preclude that the system could tip back into a previously visited set: that possibility depends on what property is establishing the equivalence classes.

Example 64. A climate change model that relates the CO_2 content of the atmosphere to global temperature may have states that are grouped together according to a shared property of those states (e.g. sea level, precipitation, glacial coverage). Due to feedback mechanisms in the system it is likely the case that these qualitative features change in punctuated equilibria thus producing equivalence classes for some states of the system. *Ex hypothesi* people can manipulate the level of CO_2 to higher or lower values. The values at which the property shifts happen may differ for the increasing and decreasing directions, but the point is that CO_2 levels could raise temperatures to the point where glaciers disappear and then later lower past the point where glaciers will form again. For some systems behavior can tip out of an equivalence class and then later tip back in. Phase transitions in condensed matter physics are another example of reversible tipping behavior. So while some have posited that tipping points are points of no return for system behavior, that turns out to be true only for certain systems and is not properly part of the definition.

Dynamics of staying, leaving, returning, and avoiding a specified set of states are covered in a separate paper on robustness (Bramson 2009). Here we continue with ways to quantify changes in what is possible for system dynamics for different states and transitions. These measures apply for any reference state or reference set, but for convenience and intuition pumping the following presentation will adopt the notation of attractor (A_i) for a reference state and set and $\mathbf{A}$ for a collection of reference states. So $\mathbf{A}$ is a collection of independent reference states and sets each of which satisfies a property while A_i may be a set of states that collectively satisfies a particular dynamical property (such as an orbit as a whole satisfies the dynamical property of an equilibrium).

Definition 617. The *energy level* of a state is the number of reference states within its reach. We write this as $E(S_i)$ and it equals

$$|\bigcup_{A_i}^{\mathbf{A}} A_i \in \mathbf{R}(S_i)|$$

Energy level quantities partition the system's states into equivalence classes.

Definition 618. The equivalence class mapping created by states' energy levels is called the system's *energy plateaus*. Each energy plateau is a set

$$\bigcup_{i,j}^{N} E(S_i) = E(S_j).$$

Definition 619. The change in energy across a transition is called an *energy precipice* or *energy drop*. We can measure the magnitude of an energy precipice in the obvious way:

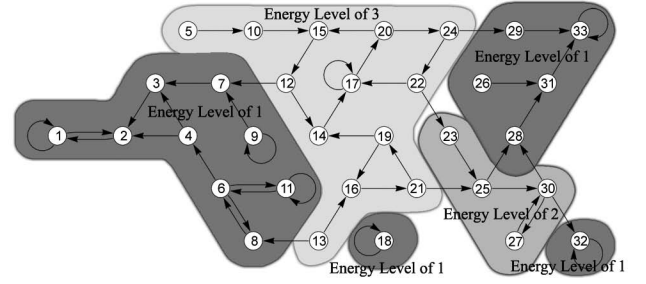$$\triangle E(\overrightarrow{S_i S_j}) = E(S_i) - E(S_j)$$

Figure 6: Energy Plateaus

Theorem 63. *An energy precipice is never negative:* $\triangle E(\overrightarrow{S_i S_j}) \geq 0$.

Proof. From definition 617 the energy level of any state S_i is the number of attractors in $\mathbf{R}(S_i)$. By theorem 53 for any $\overrightarrow{S_i S_j}$ $\mathbf{R}(S_j) \subseteq \mathbf{R}(S_i)$. This decreasing reach property implies

$$|\bigcup_{A_i}^{\mathbf{A}} A_i \in \mathbf{R}(S_j)| \leq |\bigcup_{A_i}^{\mathbf{A}} A_i \in \mathbf{R}(S_i)| \Rightarrow E(S_j) \leq E(S_i)$$

Since the end state of a transition always has a lower or equal energy level, $E(S_i) - E(S_j)$ is always greater than or equal to zero. $\square$

We now can measure the degree to which a state is likely to be the site of a tip (in a similar fashion to definition 614 of state criticality above).

Definition 620. The *tippiness* of a state is the probabilistically weighted proportional drops in energy of its immediate successors:

$$1 - \sum_{j=1}^{k} P_{ij} \left(\frac{E(S_j)}{E(S_i)} \right)$$

Theorem 64. *Tippiness ranges from zero to one.*

Proof. The lower bound occurs when all neighbors can reach the same number of reference states: $\forall k \ E(S_i) = E(S_k)$. In this case tippiness equals $1 - \sum_{j=1}^{k} P_{ij} \cdot 1 = 0$ since exit probabilities must sum to one. By theorem 63 $E(S_i) \not< E(S_k)$. When $E(S_i) > E(S_k)$ the upper bound occurs when a state has every attractor and only attractors as neighbors. We know by definition that $\sum_{j=1}^{k} P_{ik} = 1$. Attractors have an energy level of one and the energy level of S_i in this case is $|\mathbf{A}|$ so S_i's tippiness is $1 - \sum_{j=1}^{k} \frac{P_{ij}}{|\mathbf{A}|} = 1 - \frac{1}{|\mathbf{A}|}$. Thus the upper bound of S_i's tippiness goes to 1 as $|\mathbf{A}| \to \infty$. $\square$

Note that tippiness uses the ratio of energies rather than the difference; this makes tippiness a dimensionless metric and thus comparable across any state or system. Sometimes one will be more interested in minimizing the magnitude of energy drops, or avoiding states with the highest expected magnitude of energy drops, which have obvious formulations given the above definitions.

Tipping Levers As another refinement of levers from definition 61 we can apply the lever concept to tipping points to relate tips to the individual aspect changes that drive them.

Definition 621. A state has a *tipping lever* with respect to some specified set of states if a change in that aspect (or those aspects) of the state will take the system out of that set of states.

This merely combines the concept of a lever with the general concept of tipping behavior. So while all states except equilibria have levers, only perimeter states have critical levers. Identifying the tipping levers of certain sets of states is precisely what we'd like a "tipping point" analysis to reveal because it is just the aspects of a tipping point that actually change when the system tips out of a set of states. These ideas are further refined in the analysis of robustness and related concepts elsewhere (Bramson 2009).

7. Conclusions and Future Work

These tipping point measures may succeed in many cases to provide the insight necessary to understand the dynamics of complex systems, but certainly there is more work to do in refining, distinguishing, and improving these definitions. There are other dynamical properties in need of definition and improved algorithms to run the analyses on actual data. Good methodology exists as a facilitator to good science, so the first and perhaps most important extension of this project is to apply these measures to models within substantive research projects. All of these are in progress and collaboration is welcome.

As noted above, one major goal of complex systems research is to uncover the similarities across seemingly very different phenomena in a wide array of substantive domains. We can use the Markov representation to identify network *motifs* (repeated patterns in the graph structure) and establish cross-disciplinary equivalence classes of system behavior.

There are two potential non-trivial objections to the above-given probabilistic accounts of properties of system dynamics. The first is that probabilistic definitions are inadequate because we aim to understand these features as properties that systems possess rather than dynamics they *might* have. As long as the above definitions reveal useful distinctions and patterns of system behavior the project was a success, but still better (or at least different and also useful) measures may be available if built from a different formal foundation. I will, naturally, continue to pursue other and hugely different measures of system dynamics and strongly encourage input from others working on similar projects.

The other objection to the probabilistic definitions provided is that a person may insist that for many of these concepts the definition is incomplete without the causal explanation for how it comes about. Like all other statistics-like approaches these measures may be realized by many different micro-level dynamics. Some of those dynamics may not seem proper candidates for tipping behavior even if the data they generate reveals it as such from this analysis. But if this were to happen then I would consider the project a huge success. This would be similar to discovering scale-free degree distributions in many different networks from disparate research fields. Finding that common property urged researchers to pursue more deeply the phenomena and they eventually uncovered several different mechanisms by which a scale-free network may be created. Our understanding of each of those systems greatly increased because we had a common yardstick with which to measure them. The probabilistic measures presented here are not intended to replace or make unnecessary the deeper scientific analysis - they are supposed to foster it.

References

Bramson, A. 2009 "Formal Measures of Dynamical Properties: Tipping, Robustness, and Path Dependence" Forthcoming.

Clauset, A., C. Moore and M. E. J. Newman 2008. "Hierarchical structure and the prediction of missing links in networks." *Nature*. 453, 98-101.

Farrell, B. 2006. "Political Science." *The Nation*. http://www.thenation.com/doc/20060227/farrell

Gladwell, M. 2000 *The Tipping Point*. Boston, Mass: Little, Brown, and Company.

Granovetter, M. 1978 "Threshold Models of Collective Behavior" *American Journal of Sociology*. volume 83, page 1420.

Grodzins, M. 1957 "Metropolitan segregation" *Scientific American* 197(4):33-41.

Lint, J.H. Van and Wilson, R.M. 1992. *A Course in Combinatorics*. Cambridge: Cambridge University Press.

Rubineau, B. and R. Fernandez 2006 "Tipping Points: Referral Homophily and Job Segregation.'

Schelling, T. C. 1971. "Dynamic Models of Segregation" *Journal of Mathematical Sociology*. 1 (2): 143-186.

Walker, G. 2007. "The tipping point of the iceberg." in *Nature*. vol 441, no. 7065.

Wolf, E. P. 1963. "The tipping-point in racially changing neighborhoods." *Journal of the American Institute of Planners* 29(3): 217-222.

Young, H. P. 2003. "The Diffusion of Innovations in Social Networks" *The Economy as a Complex Evolving System*, vol. III, Lawrence E. Blume and Steven N. Durlauf, eds. Oxford University Press.