Structural Chaos

Conor Mayo-Wilson

June 12, 2014

Abstract

Philosophers often distinguish between parameter error and model error. Frigg et al. [2014] argue that the distinction is important because although there are methods for making predictions given parameter error and chaos, there are no methods for dealing with model error and "structural chaos." However, Frigg et al. [2014] neither define "structural chaos" nor explain the relationship between it and chaos (simpliciter). I propose a definition of "structural chaos", and I explain two new theorems that show that if a set of models contains a chaotic function, then the set is structurally chaotic. Finally, I discuss the relationship between my results and structural stability.

Climate scientists need at least two types of information to generate forecasts: (1) data about the earth's current climate and (2) a model that describes how the climate changes over time. Thus, there are at least two causes of inaccuracy in climate predictions. First, predictions might may be inaccurate because current climatic conditions are mismeasured or misestimated. Call this **initial conditions error** (ICE). Alternatively, error may arise from an inaccurate model of how the climate changes over time. Call this **structural model error** (SME).[1]

The same remarks apply to predictions about any dynamical system. If one is interested in predicting the evolution of an ecosystem over time (e.g., how population levels of various organisms change), or the behavior of markets (e.g. how prices of various commodities change), or how an epidemic will spread through a city, etc., one needs to identify both the initial conditions of the system and how the system changes over time. So there are likewise at least two sources of error in all these problems.[2]

[1] This distinction is similar to Parker [2010]'s distinction between parameter and model uncertainty.

[2] For a discussion of other sources of error in modeling, see Bradley [2012].

In a recent paper, Frigg et al. [2014] argue that the distinction between SME and ICE is crucial for both scientific practice and policy-making. They claim that, although there are methods that can generate accurate predictions in the presence of both (i) ICE and (ii) chaos, there are no known methods for doing the same with respect to (i') SME and (ii') an analogous notion of "structural chaos", which they call the "hawk-moth" effect.[3] For this reason, Frigg et al. [2014] argue that structural chaos and SME are neglected, but important topics within philosophy of science.

Although they provide an illustrative example and ample computer simulations to suggest structural chaos might be widespread, Frigg et al. [2014] do not define "structural chaos" or investigate its relationship to chaos (simpliciter).[4] This is important because there are many definitions of "chaos", and so there might be many analogous notions of "structural chaos."[5]

Frigg et al. [2014]'s arguments, therefore, raises at least three important questions for philosophers of science, applied mathematicians, and working scientists. First, for each definition of "chaos", what is the analogous concept of structural chaos? Second, what are the relationships among the various notions of chaos (simpliciter) and the analogous notions of structural chaos? Finally, what are the implications of structural chaos for prediction, control, and explanation?

This paper takes a preliminary step with respect to the first two questions. Section one describes some conditions that are used to define "chaos." I focus on topologically mixing systems, which are an important class of chaotic ones.[6] In section two, I define an analogous notion of "structural mixing" that might be used to characterize structural chaos. I then prove that, when a sufficiently rich collection of models contains a topologically mixing function, then the collection is structurally mixing in my sense.

Section three explores the relationship between my results and other

[3]Similar arguments appear in [Parker, 2011].

[4]Frigg et al. [2014] do formally define what they call "closeness to goodness fit." This definition is analogous the the definition of sensitivity to initial conditions, which is generally considered to be a necessary but insufficient condition for chaos. See section one below. At points, they implicitly suggest that structural *in*stability might be the structural analog to chaos. This suggestion is criticized in the section three.

[5]For discussions of definitions of chaos, see [Batterman, 1993] and [Werndl, 2009].

[6]According to Devaney et al. [1989]'s widely-cited definition, a system is chaotic if it satisfies three conditions: (i) it is sensitive to initial conditions; (ii) it is topologically transitive, and (iii) its periodic points are dense in state space. Topological mixing systems are topologically transitive, and under very general conditions, they are also sensitive to initial conditions. Thus, they satisfy two of the three properties that are widely used to define "chaos."

potential characterizations of structural chaos. In particular, I argue that definitions of "structural instability", which are often informally motivated in ways analogous to definitions of chaos, are not clearly analogous to notions of chaos simpliciter. The final section discusses the philosophical importance of my results and answers to the above three questions.

1 Chaos

Popular writings often describe chaos via an appeal to Lorenz's metaphor of the "butterfly effect". Lorenz famously asked whether the flapping of a butterfly's wings in Brazil could cause a thunderstorm in Texas. In general, a chaotic system is often described as one in which small changes (e.g. a butterfly flapping its wing) in the initial conditions of a system can create large changes in its behavior (e.g., storm patterns).

This informal gloss captures only one aspect of standard definitions of "chaos", however. To give more precise characterizations, it is necessary to introduce some definitions. A discrete-time **dynamical system** is a triple $\langle X, d, \varphi \rangle$ where (i) $\langle X, d \rangle$ is a metric space called the **state space**, and (ii) $\varphi : X \to X$ is a **time-evolution** function.[7] For the remainder of the paper, I use the phrases "model", "dynamical function" and "time-evolution function" interchangeably, though of course I recognize not all models in science are time-evolution functions.

For example, a dynamical system might describe the motion of a particle in space. In this case, X is be three-dimensional space; d represents a function specifying the distance between points in three-dimensional space, and φ is a function describing how a particle moves over time. Or X might be the set of vectors specifying the temperature, pressure, and wind velocities at different places in the atmosphere; d would represent how similar two descriptions of the earth's climate are, and φ would represent how the climate changes over time.

How can one use the definition of a dynamical system to capture the notion of sensitivity to initial conditions? Let Δ be a number representing a large distance between states. What counts as "large" can depend upon the state space and one's interests. Say a dynamical system's behavior is **sensitive to initial conditions** to degree Δ if for every state $x \in X$ and every arbitrarily small distance $\epsilon > 0$, there exists a state y within distance ϵ of x and a natural number N such that $d(\varphi^N(x), \varphi^N(y)) > \Delta$. Here, $\varphi^N(x)$

[7] Note that, for simplicity, I assume that the time evolution function φ is constant over time. Not all discrete dynamical systems have this property.

represents the state of the system after N stages if its initial conditions were x. Informally, a system exhibits sensitivity to initial conditions if no matter the true initial state x, there is an arbitrarily close state y such that, if y had been the initial state, the future would have been radically different.

This mathematical definition is the natural way of capturing the above informal description of chaos above, but there are many time-evolution functions that are sensitive to initial conditions in the above sense and yet are hardly "chaotic" in any sense of the word. Consider, for example, the function $f(x) = 2x$ on the state space consisting of all real numbers. Then f is sensitive to initial conditions because if two numbers x and y differ by even the smallest amount, then the result of multiplying them by two repeatedly will cause them to drift apart. That is, $|f^n(x) - f^n(y)| = 2^n|x - y|$ becomes arbitrarily large as n grows. So f is sensitive to initial conditions, but f does not exhibit "chaotic" behavior in the least.

What other conditions might one add in order to characterize "chaos"? It turns out there is no wide agreement, and that several different definitions of chaos are common.[8] Because my aim is to show how three types of questions might be answered (see above), I will not defend a particular analysis of chaos. Rather, I will simply show how to answer the three questions with respect to the concept of "topologically mixing", which plays an important in characterizing chaos (see footnote 5).

A time-evolution function φ is called **topologically mixing** if for any pair of non-empty open sets U and V, there exists a number $N > 1$ such that

$$\varphi^n(U) \cap V \neq \emptyset.$$

for all $n \geq N$. In order to reduce the amount of technical jargon, I will say φ is **chaotic** if it is topologically mixing.

For the reader unfamiliar with topology, ignore the phrase "open set" for now. Just think of U and V as representing sections of state space. If the system begins in some state in U, then the expression $\varphi^n(U)$ represents all possible future states after n many steps of time. For example, suppose the dynamical system describes the movement of a gas molecule in a room. Further, assume that U represents the upper-left quarter of the room and that V represents the lower-right hand corner. Then $\varphi^n(U)$ represents the

[8]For what it's worth, I agree with Werndl [2009] that the vast majority of systems that are agreed to be chaotic are strongly mixing in the sense of ergodic theory. Moreover, I agree with [Berkovitz et al., 2006] that, because strong mixing is one among several logically related concepts of probabilistic independence in the ergodic hierarchy, it is probably most productive to think of chaos as coming in degrees, where different degrees may have different implications for prediction, explanation, and control.

possible positions of the gas molecule after n many units of time if the gas particle started in the upper-left quarter of the room. The above equation says that there is some time in the future such that, from that point onward, there is always a position in the upper-left corner of the room (U) such that, if the gas particle had started in that position, then it would end up in the lower-right quarter of the room (V). A time-evolution function is chaotic if this holds for any regions of state space, which is to say that (in the example) a gas particle that starts in one section of the room can end up in any other section of the room after a sufficiently large period of time.

If topological mixing is taken to be a characteristic of chaotic systems, would would it mean to say that SMEs can lead to "structural chaos"? This is the topic of the next section.

2 Structural Chaos

A dynamical system is chaotic if, when the time-evolution function is held fixed, similar initial conditions can have any future. Analogously, a set of dynamics should be called "structural chaotic" if, when the initial conditions are held fixed, similar time-evolution functions can produce any future. See figure below. To rigorously define "structural chaos", therefore, one needs a metric to quantify how "close" two time-evolution functions are.

Let X^X represent all time-evolution functions for a system with state space X. Depending upon one's interests, there are different appropriate metric quantifying the distance between models (i.e. time-evolution functions). However, clearly there is some relationship between (1) the distance between two models and (2) the distances between their predicted future states after one unit of time. If two models entail that a system, starting in the same initial position, will be in radically different places in a short amount of time, then the models are substantially different.

One demanding notion of closeness requires that two models are close precisely if their values are *always* close. In other words, the distance between two time-evolution functions is the maximum/supremum distance between the models after one unit of time, where the maximum is taken over all possible starting states. In symbols, define:

$$D(\varphi, \psi) = \sup_{x \in X} d(\varphi(x), \psi(x)).$$

Henceforth, I assume that D quantifies the distance between two time-evolution functions, but the results below hold for a variety of metrics.

"Structural mixing" should capture the idea that similar models can produce different trajectories through the state space given the same initial conditions. To make this idea rigorous, I introduce some notation. For any $\epsilon > 0$, let $B_\epsilon(\varphi)$ denote all models within distance ϵ of φ. Next, for any natural number $n \in \mathbb{N}$ and any point $x \in X$, define a map $f_{x,n} : \mathcal{P}(X^X) \to \mathcal{P}(X)$ as follows:

$$f_{x,n}(\Phi) = \{\varphi^n(x) : \varphi \in \Phi\}$$

where $\mathcal{P}(X)$ is the power set of X, i.e., the set of all subsets of X. In other words, $f_{x,n}$ maps a set of time-evolution functions to the set of points they reach after n stages if they are initialized to start at x.

Given a set of time-evolution functions $\Phi \subseteq X^X$ and a particular model $\varphi \in \Phi$, say that Φ is **structurally mixing at** φ if for all $x \in X$, all $\epsilon > 0$ and all non-empty open sets $V \subseteq X$, there is some $N \in \mathbb{N}$ such that

$$f_{x,n}(B_\epsilon(\varphi) \cap \Phi) \cap V \neq \emptyset$$

for all $n \geq N$. In other words, small differences between the estimated model and the true one can lead to divergent predictions *even if one correctly identifies the initial condition*. To reduce jargon, I sometimes say a set of models is **structurally chaotic** at φ if it is structurally mixing.

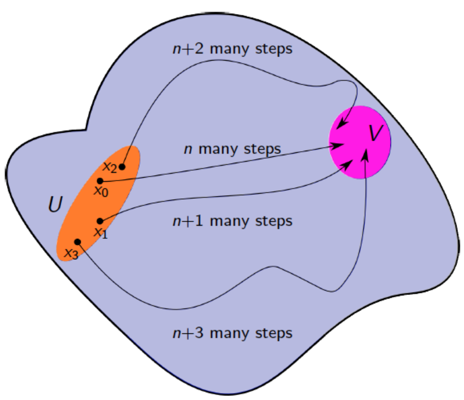

Topological Mixing

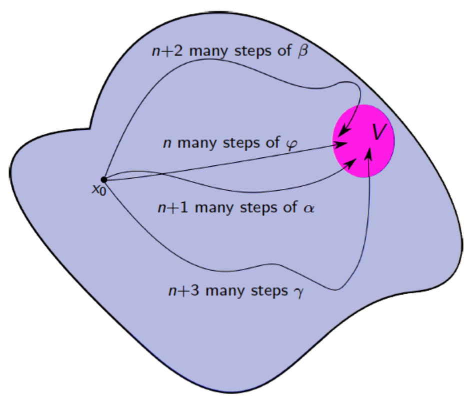

Structural Mixing

The concept of structural mixing is the obvious analog of the definition of topological mixing in the previous section. Clearly, different definitions of chaos will generalize to different definitions of structural chaos. Nonetheless, this example suggests a new research program, which consists of three questions. First, for each definition of "chaotic system", what is the analogous concept of structural chaos? Second, what is the relationship between the various notions of chaos (simpliciter) and the analogous notions of structural chaos? Finally, what are the implications of structural chaos for prediction, control, and explanation?

Given my definition of structural chaos, the second question can be given a precise answer:[9]

Theorem 1 *Suppose φ is continuous and topologically mixing. If X has no isolated points, then X^X is structurally mixing at φ.*

That is, the set of possible time-evolution functions is structurally chaotic if it contains a chaotic model. One might object that this theorem is very weak. According to the theorem, one should worry about structurally chaos if *every* time-evolution function were a plausible description of the dynamics of the system. However, in practice, the set of plausible models is much narrower given existing data, domain knowledge, physical constraints, and so on. For example, if it were $40°C$ in Demascus today, then it would be bizarre if it snowed tomorrow. However, one possible time-evolution function for Demascus' weather will entail that a $40°C$ day will be followed by a snowy day. Thus, one might object that if the class of models is restricted to realistic time-evolution functions, then structural chaos will be rarer.

However, the proof of the above theorem actually shows something much stronger. It shows that, if the true time-evolution function is chaotic and the set of possible time-evolution functions contains all models that are empirically indistinguishable from the true one, then structural chaos will arise. To explain why, I introduce some definitions.

Data sets are always finite. So let $F = \{x_0, \ldots, x_n\}$ by a finite set of states, which represents the observed history of the system so far. Let $\epsilon > 0$ be a small number representing the precision of one's measurement devices. Say two models are **ϵF-indistinguishable** if (1) the values of time-evolution functions are equal for all but finitely many states outside F and (2) the two models are no more than ϵ apart according to D.

Two models are ϵF-indistinguishable if they are, in a very strong sense, indistinguishable given all available empirical data. Why? The first clause entails that the two models are equal on all observed data points, and so there is no way that past data alone can distinguish between them. If two models differ *anywhere*, however, then there are logically possible experiments that can distinguish between them. Namely, if controlled experiments are financially, pragmatically and ethically feasible (which they often

[9]See appendix for a proof.

are not), one can initialize the system to one of the states at which the two models differ and observe the results.

This is where the second clause kicks in. Suppose scientists' measuring instruments and statistical techniques cannot guarantee estimates of the observed states with accuracy better than $\epsilon > 0$. If two models are ϵF-indistiguishable, then second clause guarantees that no information about the current or next state of the system is sufficient to distinguish the models. One might object that small measurement errors are detectable in the long run, especially if the model is chaotic. However, if the true dynamics are continuous and ϵ is sufficiently small, then the second clause entails that no experiment of a feasible length (i.e. time) will distinguish between it and an ϵF-indistinguishable model.

The previous discussion motivates the following definition. Let F denote the finite set of observed states. Say a set Φ of time-evolution functions is **closed under empirical-indistinguishability** if there exists some $\epsilon > 0$ such that if $\varphi \in \Phi$ and ψ is ϵF-indistinguishable from φ, then $\psi \in \Phi$. The above argument is intended to show that, if scientists are strict empiricists, then the set of models that they consider possible ought to be closed under empirical indistinguishability. Theorem 1 is a special case of the following stronger result.

Theorem 2 *Suppose φ is continuous and chaotic. Let Φ be a set of time-evolution functions containing φ. If X has no isolated points and Φ is closed under empirical indistinguishability, then Φ is structurally chaotic at φ.*

3 Structural Stability: Conclusions and Future Research

Readers familiar with chaos theory may find the previous theorem surprising. On one hand, my definition of "structural chaos" seems to formalize the idea that small errors in identifying the model can lead to divergent future behavior. On the other hand, many of the time-evolution functions that lead to "structural chaos" (according to my definition) are *structurally stable* in one or more senses.[10] This is counter-intuitive because structural stability

[10]Suppose $f : A \to A$ and $g : B \to B$ are functions on topological spaces. Then f and g are said to be *topologically conjugate* if there is a homeomorphism $h : A \to B$ such that $g \circ h = h \circ f$. A function $f : A \to A$ is C^r *structurally stable* if there is some $\epsilon > 0$ such that every function within distance ϵ of f in the C^r metric is topologically conjugate to f. C^r structural stability is perhaps the most common definition, but other definitions have a similar logical form, which is discussed in the body of the paper.

is intended to formalize the idea that small changes to the model do not result in large differences in the model's trajectory.

It is best to begin with an example to understand the tension. A paradigmatic chaotic function is the *logistic map* $F_r(x) = rx(1-x)$, where r is greater than about 3.57. It is known that the logistic map is C^2-structurally stable when $r > 4$, and it is structurally stable on an open dense set of values of r between 0 and 4. For this reason, some chaos theorists might claim that small changes to the logistic map will not result in divergent future behavior. However, the logistic map (for $r = 4$) is precisely the example that Frigg et al. [2014] use to demonstrate the impacts of structural chaos. Moreover, if Φ is a set of models that contains the logistic map and is closed under empirical-indistinguishability, then Theorem 2 entails that Φ is structurally chaotic at the logistic map, as the logistic map is topologically mixing. So Frigg et al. [2014]'s and my results seem to be in tension with facts about structural stability.

One possible reason for the tension is that definitions of structural stability almost always assume that the set of models under investigation are well-behaved, in the sense that models are differentiable (perhaps several times) and hence, continuous. In contrast, in order to demonstrate the existence of "structural chaos" in computer simulations, Frigg et al. [2014] simulate discretized functions that are, by necessity, discontinuous. Moreover, if a set of models is closed under empirical indistinguishability in my sense, it will contain discontinuous functions and other "poorly behaved" models. Some may see this as a deficiency in Frigg's and my arguments. Continuity and differentiability are mathematically convenient assumptions, and Ockham's razor or other metaphysical arguments might lead one to accept that the dynamics of real physical systems are continuous. Nonetheless, convenience and simplicity are extra-empirical considerations; a finite sequence of observed states may be consistent with assuming the continuity of the system's time-evolution function, but it does not require doing so. Furthermore, many metaphysical arguments for continuity do not obviously extend to showing that a function is twice differentiable.

However, I will not defend the thesis that physical laws might be discontinuous or non-differentiable. Rather, I discuss the relation between structural chaos (in my sense) and various notions of structural stability in order to illustrate a broader point. Mathematicians, scientists, and philosophers have yet to investigate whether plausible structural analogs of "chaos" are actually in tension with definitions of structural stability. My results show that there may, in fact, be no direct logical inconsistency, and that inconsistency may only arise when additional, substantive assumptions (e.g. continuity or differentiability) about the dynamics of the system are introduced.

There are two further reasons to question whether standard definitions of "structural instability" are really the appropriate dynamical analogs of chaos. To understand the two reasons, it is not necessary to review all existing definitions of structural stability. Rather, it suffices to describe their common logical form [Pugh and Peixoto, 2008]. Namely, given some equivalence relation R (e.g., topological conjugacy) over functions, one says a function f is structurally stable if all "close" functions (under some metric) are R-equivalent to f. Why are definitions of this form not analogous to characterizations of chaos (simpliciter)?

First, the concepts employed to define structural stability are disjoint from those used to define chaos. For example, definitions of structural stability typically use the notions of homeomorphism and diffeomorphism, whereas definitions of chaos employ notions like sensitivity to initial conditions, topological transitivity, density, etc. Of course, some difference in definitions is unavoidable, as structural stability is about small changes in time-evolution functions, whereas chaos is about small changes in states.

Nonetheless, if Werndl [2009] and Berkovitz et al. [2006] are correct, then probability is a key concept in characterizing degrees of chaos. In contrast, none of the definitions of structural stability employ probability at all. This is surprising, given that probability (and in particular, probabilistic independence) is perhaps the most widely-employed tool used to characterize uncertainty, noise, and (expected) error. The fact that probability is not used in definitions of structural stability, therefore, raises serious questions about the importance of such definitions for discussions of prediction, control, and explanation.[11]

Second, time plays different roles in definitions of chaos and structural stability respectively. Definitions of chaos typically contain a clause – like the definition of topological mixing – that places constraints on the distant future of the system. For example, in many chaotic systems, nearby initial conditions may have similar trajectories for a long period of time, but their trajectories may diverge suddenly and radically in the distant future. The potential for such sudden divergence is what renders long-term predictions problematic. In contrast, to my knowledge, all but one of the equivalence relations used to define structural stability constrain *only one time step* in the evolution of a dynamical system, and the one exception is typically only

[11] The reader will note that my definition of structural mixing likewise does not employ the use of probability. It turns out that the standard notion of topological mixing is closely related the ergodic (and hence, probabilistic) concept of strong mixing. I conjecture an analogous relationship will hold in the structural case, but this remains to be shown.

applied to dynamical systems that are described by differential equations.

These two reasons do not provide conclusive evidence that the mathematically rich research on structural stability is, at the end of the day, unimportant for empirical science. Rather, they suggest two more questions to add to the list at the outset of the paper: what are the relationships among various definitions of chaos and structural stability? And what is the importance of the various notions of structural stability for prediction, control, and explanation?

4 Conclusions and Philosophical Upshots

Section one outlined a broad research program, which consisted of three questions. Section two provided a brief example of how one might go about answering two of three questions. In particular, I defined a notion of "structural mixing" that is analogous to the standard notion of "topological mixing", and I proved a theorem relating the two concepts. I now conclude by discussing philosophical significance of this research program.

To see why this seemingly technical series of questions has broad philosophical importance, replace every occurrence of the phrase "time-evolution function" with the word "regularity" in the above discussion of structural chaos and in the two theorems. Doing so reveals that the main result roughly asserts that there are many "similar" regularities that (i) produce widely different future behavior and (ii) are compatible with the observed past. That's just an instance of the problem of induction. So investigating structural chaos amounts to investigating (in a mathematically precise setting) a (the?) central problem of epistemology and philosophy of science.

With this in mind, it is now easy to see why answers to each of the three questions are philosophically important. Question one asks, "For each definition of "chaotic system", what is the analogous concept of structural chaos?" Because there are different "degrees" of chaos [Berkovitz et al., 2006], an answer to question one would characterize differing "degrees" of problem of induction.[12] That is, an answer to the first question would allow one to characterize inductive problems in terms of their difficulty.

Question two asks, "what are the relationships among the various notions of chaos (simpliciter) and the analogous notions of structural chaos?" To see

[12] Kelly [1996] contains a sophisticated description of a hierarchy of "problems" of induction. I am skeptical there is any relationship between Kelly's hierarchy and that which would arise from pursuing the first question here. So this project would provide an independent, orthogonal way of characterizing inductive difficulty.

why this question is important, it is useful to consider one reason why chaotic systems are so interesting. The classic problem of induction shows that past observations are insufficient to identify a dynamical system's time-evolution function, and hence, there are many regularities that (a) are compatible with past observations and (b) predict radically different futures. The existence of chaos entails that predicting or manipulating a dynamical system's behavior might be difficult *even if the exact dynamics of the system are known.* Hence, an answer to question two provides a bridge between research on the classical problem of induction and new research in chaos theory, which respectively identify different sources of difficulty for prediction and manipulation.

Finally, question three asks, "what are the implications of structural chaos for prediction, control, and explanation?" The importance of this question is self-explanatory: prediction, control, and explanation are three central goals of science, and so an answer to question three amounts to an answer to the question, "Why is structural chaos important?"

A Proofs

Lemma 1 *Let X be any metric space, $U \subseteq X$ an open set and $F \subseteq X$ be finite. Then $U \setminus F$ is open. If X be has no isolated points, $U \setminus F$ is non-empty.*

Theorem 2 *Suppose φ is continuous and topologically mixing. Suppose that $\varphi \in \Phi$ and that Φ is closed under F-indistinguishability for some finite $F \subseteq X$. If X has no isolated points, then Φ is structurally mixing at φ.*

Proof: Let $x_0 \in X$. It must be shown that for all $\epsilon > 0$ and all non-empty open sets $V \subseteq X$, there is some $N \in \mathbb{N}$ such that

$$f_{x_0,n}(B_\epsilon(\varphi) \cap \Phi) \cap V \neq \emptyset \text{ for all } n \geq N$$

Call this condition $\dagger(\epsilon, V, N)$. Let $\epsilon > 0$ and $V \subseteq X$ be an open set.

Define $x_j = \varphi^j(x_0)$ for all natural numbers j, and let $M = |F| + 1$. Because Φ is closed under F-indistiguishability, there is $\beta > 0$ such that if (a) φ and ψ agree everywhere on all but finitely many elements of $X \setminus F$ and (b) $D(\varphi, \psi) < \beta$, then $\psi \in \Phi$. As φ is continuous and F is finite, it follows that for all $k \leq M$ there is $\delta_k > 0$ such that

$$B_{\delta_k}(x_k) \cap F = \begin{cases} \{x_k\} & \text{if } x_k \in F \\ \emptyset & \text{otherwise.} \end{cases}$$

and

$$y \in B_{\delta_k}(x_k) \Rightarrow d(\varphi(y), \varphi(x_k)) < \{\epsilon, \beta\}$$

Note here I am using $B_\gamma(z)$ to refer to the γ-ball around $z \in X$ with respect to the metric d, in the same way that I have used $B_\gamma(\varphi)$ to refer to the γ-ball around φ with respect to D.

Let $\delta = \min\{\delta_k : k \leq M\}$. Because φ is topologically mixing, for each $k \leq M$ there is $N_k \in \mathbb{N}$ such that for all $n \geq N_k$:

$$\varphi^n(B_\delta(x_k)) \cap V \neq \emptyset$$

Let $N_* = M + \max\{N_k : k \leq M\}$. I claim that $\dagger(\epsilon, V, N_*)$. Let $n \geq N_*$. It is necessary to find a function $\psi \in B_\epsilon(\varphi) \cap \Phi$ such that $\psi^n(x_0) \in V$. If $\varphi^n(x_0) \in V$, then we're done. So assume $\varphi^n(x_0) \notin V$.

Because $M > |F|$, there is $k \leq M$ such that $x_k \notin F$. Notice

$$n - k \geq N_* - M \geq \max\{N_j : j \leq M\} \geq N_k.$$

Hence, by choices of δ and N_*, there is $y \in B_\delta(x_k)$ such that $\varphi^{n-k}(y) \in V$. Note $y \neq x_k$ because $\varphi^{n-k}(x_k) = \varphi^n(x_0) \notin V$. I claim that y may be chosen so that $\varphi^j(y) \neq x_k$ for all $j \leq n - k$.

Why? Suppose for the sake of contradiction that for all $y \in B_\delta(x_k)$, there is some $j \leq (n-k)$ such that $\varphi^j(y) = x_k$. In particular, there is $j_0 \leq (n-k)$ such that $\varphi^{j_0}(x_k) = x_k$. Thus, for all $m \geq (n-k)$ and all $y \in B_\delta(x_k)$:

$$\varphi^m(y) \in \{x_k, \varphi(x_k), \ldots, \varphi^{j_0-1}(x_k)\}.$$

Let $T = X \setminus \{x_k, \varphi(x_k), \ldots, \varphi^{j_0-1}(x_k)\}$. Then T is non-empty and open by the lemma. However, by the above reasoning, $\varphi^m(B_\delta(x_k)) \cap T = \emptyset$ for all $m \geq (n-k)$. So φ is not topologically mixing, contradicting assumption.

It has been shown that $y \in B_\delta(x_k)$ may be chosen so that $\varphi^j(y) \neq x_k$ for all $j \leq (n-k)$. Define $\psi : X \to X$ as follows:

$$\psi(z) = \begin{cases} \varphi(y) \text{ if } z = x_k \\ \varphi(z) \text{ otherwise.} \end{cases}$$

Note $D(\varphi, \psi) = d(\varphi(x_k), \varphi(y))$. By continuity of φ, it follows that $d(\varphi(x_k), \varphi(y)) \leq \min\{\beta, \epsilon\}$. Hence, $\psi \in B_\epsilon(\varphi)$. Because ψ is equal to φ everywhere except $x_k \notin F$, it follows that ψ is βF-indistinguishable from φ. As Φ is closed under βF-indistinguishability, $\psi \in \Phi$.

Finally, $\psi^n(x) = \varphi^{n-k}(y) \in V$ because $\varphi^j(y) \neq x_k$ for all $0 \leq j \leq n-k$.

$\square$

References

Robert W. Batterman. Defining chaos. *Philosophy of Science*, 60(1):43—66, March 1993.

Joseph Berkovitz, Roman Frigg, and Fred Kronz. The ergodic hierarchy, randomness and hamiltonian chaos. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 37 (4):661–691, 2006.

Seamus Bradley. Scientific uncertainty: A user's guide. *Grantham Institute on Climate Change*, Discussion Paper 56, 2012. URL `http://philpapers.org/rec/BRASUA`.

Robert L. Devaney, Luke Devaney, and Luke Devaney. *An introduction to chaotic dynamical systems*, volume 6. Addison-Wesley Reading, 1989.

Roman Frigg, Seamus Bradley, Hailiang Du, and Leonard A. Smith. Laplaces demon and the adventures of his apprentices. *Philosophy of Science*, 81(1):31—59, 2014.

Kevin T Kelly. *The logic of reliable inquiry*. Oxford University Press, New York, 1996.

Wendy S. Parker. Predicting weather and climate: Uncertainty, ensembles and probability. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 41(3):263–272, 2010.

Wendy S. Parker. When climate models agree: The significance of robust model predictions. *Philosophy of Science*, 78(4):579–600, 2011.

Charles Pugh and Mauricio Peixoto. Structural stability. *Scholarpedia*, 3(9), 2008. ISSN 1941-6016. doi: 10.4249/scholarpedia.4008. URL `http://www.scholarpedia.org/article/Structural_stability`.

Charlotte Werndl. What are the new implications of chaos for unpredictability? *The British Journal for the Philosophy of Science*, 60(1):195–220, 2009.