

On two mathematical definitions of observational
equivalence:
Manifest isomorphism and ε -congruence reconsidered

Christopher Belanger
chris.belanger@utoronto.ca

Abstract

In this article I examine two mathematical definitions of observational equivalence, one proposed by Charlotte Werndl and based on *manifest isomorphism*, and the other based on Ornstein and Weiss's *ε -congruence*. I argue, for two related reasons, that neither can function as a *purely* mathematical definition of observational equivalence. First, each definition permits of counterexamples; second, overcoming these counterexamples will introduce non-mathematical premises about the systems in question. Accordingly, the prospects for a broadly applicable and purely mathematical definition of observational equivalence are unpromising. Despite this critique, I suggest that Werndl's proposals are valuable because they clarify the distinction between provable and unprovable elements in arguments for observational equivalence. *Keywords:* Determinism, Indeterminism, Stochastic processes, Dynamical systems, Observational equivalence, ε -congruence

1. Introduction

1.1. Background and Outline

Observational equivalence is a fraught topic in the philosophy of science, and there is little agreement as to when—if at all—it holds between scientific models. If the observational equivalence of distinct models could be reduced to a provable mathematical relation, this would be an important development in the epistemology of science. Charlotte Werndl has recently argued that two such relations, to be introduced shortly, could be useful in mathematical definitions of observational equivalence (Werndl, 2009a). In this paper I will consider whether these two relations can be considered *purely* mathematical definitions of observational equivalence, and argue that this is unlikely.

The definitions to be considered apply to deterministic dynamical systems and to stochastic processes, and so I will begin with a brief introduction to the relevant mathematics. Next, Werndl’s two proposed definitions will be considered in turn. The first is based on manifest isomorphism and is, to my knowledge, original to Werndl (2009a). The second is based on ε -congruence, which was introduced in Ornstein and Weiss (1991). I will present two main arguments for the position that neither of these relations is an acceptable condition for observational equivalence. First, I will present counterexamples which demonstrate that each definition can hold between models which are easily distinguishable, and so each proposal is insufficient as stated. Second, I will consider whether these counterexamples could be avoided by making the definitions more restrictive. I will argue that, while this is the case, these restrictions will take the form of physical hypotheses about the systems in question. This undermines the idea of a *purely* math-

emathical definition of observational equivalence, and limits any application of such modified definitions to a given theoretical context. Accordingly, I suggest that the prospects for a non-contextual, provable, and purely mathematical definition of observational equivalence are dim. In closing, I will suggest that a strength of Werndl’s analysis is that it can deepen our understanding of observational equivalence by clarifying this distinction between the provable and unprovable elements of our scientific judgements.

1.2. Deterministic Dynamical Systems

In the broadest terms, dynamical systems are mathematical representations of how values change over time.¹ Often these values are given physical interpretations, and the dynamical system can be said to model some part of the real world. The framework here is that of ergodic theory, and as such a measure-preserving dynamical system is a quadruple (M, Σ_M, μ, T) . M is a non-empty set, called the phase space, whose members represent the possible states of the system. Σ_M is a σ -algebra on M , which is a set of subsets of M that must be nonempty, closed under complementation, closed under countable unions, and must contain \emptyset and M .

μ is the *probability measure*, a function whose domain is the elements of Σ_M and whose codomain is $[0, 1]$. Intuitively, μ gives the “size” of the regions of M contained in Σ_M , and we make the controversial but common assumption here that $\mu(m)$ is the probability that the system will be found in a region m of M . μ is defined such that $\mu(M) = 1$, and is closed under countable additions. Any set assigned measure zero is called a null set, but

¹For an accessible introduction to measurable dynamical systems, see Silva (2008).

note that a set of measure zero need not be small in any intuitive sense.

The transformation T moves points and regions of M to other points and regions of M , and is interpreted as the dynamical system's evolution operator. The present paper is concerned with discrete-time dynamical systems, and so if a system's state is initially $m \in M$, then $T(m)$ is the new state of the system after one time step. This process can be iterated, and $T^n(m)$ denotes the state of the system after n time steps. Repeated iteration of T generates a sequence of points called an *orbit* or *trajectory* of the system. The types of dynamical systems considered here are both invertible and measure-preserving, and so $\mu(T^{-1}(m)) = \mu(T(m))$ for all $m \in \Sigma_M$.

A deterministic system (M, Σ_M, μ, T) is called *Bernoulli* if it meets two conditions. First, and most relevantly for our purposes here, there must exist a partition α of M such that $\mu(T^n \alpha_i \cap T^m \alpha_j) = \mu(T^n \alpha_i) \mu(T^m \alpha_j)$, where $-\infty < m, n < \infty$ and α_i and α_j range over all the atoms of α . Second, the $T^i \alpha$ must generate the full σ -algebra of M (Ornstein, 1970, 337). Setting aside this second condition, which does not bear directly on the arguments of this paper, note that the first condition states that the past and future states of a Bernoulli system exhibit a kind of statistical independence. In a well-defined sense Bernoulli systems are the most unpredictable of all deterministic systems (Berkovitz et al., 2006, 677). The KS-entropy of a Bernoulli system is defined as the following sum over its transition probabilities p_i :

$$-\sum_{i=1}^n p_i \ln p_i \tag{1}$$

Ornstein famously demonstrated that Bernoulli systems are isomorphic if and only if they have the same KS-entropy, and this result will be useful in a later section (Ornstein, 1970).

1.3. Stochastic Processes

Stochastic processes have irreducibly probabilistic state transitions, and so their evolutions cannot be predicted with certainty even given precise knowledge of their dynamics and current state. Despite stochasticity’s intuitive connection to ‘chancy’ or indeterministic processes, the fact that a physical system can be modelled by a stochastic process must not be taken as proof that it is indeterministic in some ontological sense.² Epistemic limitations, for example, will sometimes require us to use stochastic models even when we believe the systems under study are deterministic.

To define a stochastic process, let $\bar{M} = \{m_1, m_2, \dots, m_n\}$ be a set of states identified as possible outcomes. In analogy to the way we constructed a deterministic system, let Ω be a phase space, Σ_Ω be a σ -algebra, and ν be a probability measure. Intuitively, Ω represents all possible sequences of outcomes, and each $\omega \in \Omega$ represents one possible sequence of outcomes. Let us also define a set of functions Z_t from Ω to \bar{M} , which we call random variables. Z_t is interpreted as the outcome of ω at time t . Since we are concerned with discrete processes, the time index will range over the integers. Our focus here is on *stationary* stochastic processes, for which the transition probabilities are constant through time. Putting these pieces together, a stochastic process is a one-parameter family of such random variables:

$$\{Z_t; t \in \mathbb{Z}\} \tag{2}$$

²This connection is often made in the physics literature, eg: “Observe, however, that the system is still considered deterministic: only the model becomes stochastic ... this is quite different from the common approach that assumes the system must be stochastic too” (Judd and Smith, 2004, 232).

where, again, the values given by Z_0 , Z_1 , and so on represent outcomes of the process at times $t = 0, 1, \dots$. As long as the transition probability from at least one state to another has a probability that is strictly between 0 and 1 we call the stochastic process *non-trivial* (Doob, 1953, 46–47).

There are several different types of stochastic behaviour, and in this paper we will be concerned particularly with *Bernoulli* and *Markov* processes. A Markov process’ transition probabilities depend only on its current state, and on none of its past states. Bernoulli processes, on the other hand, have ‘no memory,’ and their transition probabilities are state independent. As they are stochastic, Bernoulli processes are conceptually distinct from the deterministic Bernoulli systems defined in the previous section.

1.4. *Finite-valued observation function*

In practice, our observations are always limited to some finite precision, and we can represent this mathematically with an *observation function* Φ . Let the observation space M_O represent the values we can actually observe in practice. M_O is often generated by applying a partition to the full phase space M , which divides M up into non-overlapping regions called *atoms* (see figure 1). Each atom of the partition α_i is associated with an observable value, and the observation function Φ takes each point of the phase space M to the value of its corresponding atom. This is often referred to as a ‘coarse-graining’ of the phase space. The size of the atoms will depend on factors like the resolution of our instruments, and upgrading our instruments would correspond to using a ‘finer-grained’ partition, but as long as we impose a finite partition we can never know the system’s precise state. Here we assume M_O has finitely many elements, and Φ is called a *finite-valued observation*




Figure 1: An example of a finite partition dividing M into the five atoms α_1 through α_5 . An observation function Φ would then associate an observable value with each α_i .

function.

1.5. Deterministic Representation of a Stochastic Process

The last piece of technical apparatus needed is called the *deterministic representation of a stochastic process*. This is a method for ‘replacing’ any stochastic process $\{Z_t; t \in \mathbb{Z}\}$ from $(\Omega, \Sigma_\Omega, \nu)$ to $(\bar{M}, \Sigma_{\bar{M}})$ with a deterministic system. Recall that each $\omega \in \Omega$ represents an infinite sequence of possible outcomes of the stochastic process, which is called the *realization* of ω . The idea is to set up a deterministic left-shift system whose phase space consists of all possible realizations, and gives the same probabilities as the stochastic process. Let the phase space M of the deterministic system be the set of all bi-infinite sequences $(\dots m_{-1} m_0 m_1 \dots)$ with each m_i a member of the outcome space \bar{M} . Let the transformation $T : M \rightarrow M$ be the left shift, which bumps each element over one place by moving m_i to m_{i-1} . Only the 0th element m_0 can be observed, and this is represented mathematically by observation function $\Phi : M \rightarrow \bar{M}, \Phi(m) = m_0$. As we keep applying the left shift T , different elements of the sequence will be moved to m_0 and become visible when we observe the system through Φ . The finite realizations of $\{Z_t; t \in \mathbb{Z}\}$ are cylinder sets, and the probability assigned to these realizations forms a pre-

measure which may be extended to the measure μ on Σ_M . $(M, \Sigma_M, \mu, T, \Phi)$ as so constructed is a deterministic system which reproduces the given realizations and probabilities of the stochastic process $\{Z_t; t \in \mathbb{Z}\}$, and is called its deterministic representation (Werndl, 2009a, 236).

2. Two Mathematical Definitions of Observational Equivalence

2.1. Manifest Isomorphism

If two probabilistic models assigned (nearly) the same probabilities to the same outcomes, we might want to call them observationally equivalent; conversely, if the models assigned very different probabilities to the same outcomes, it would be unintuitive to call them observationally equivalent. Charlotte Werndl has suggested that in some cases manifest isomorphism, a special case of general measure-theoretic isomorphism, could provide a rigorous justification of this kind of probabilistic, evidence-based observational equivalence (Werndl, 2009a, 234). After outlining manifest isomorphism itself, I will present counterexamples to show that models can be manifestly isomorphic but not observationally equivalent, and observationally equivalent without being manifestly isomorphic. A strengthened version of manifest isomorphism might avoid these difficulties, but I will argue that this strength comes at the cost of introducing contextual and defeasible assumptions. Such a strengthened proposal could help us gain a better understanding of observational equivalence, but would not be a purely mathematical and provable relation.

Two measure-preserving systems $(M_1, \Sigma_{M_1}, \mu_1, T_1)$ and $(M_2, \Sigma_{M_2}, \mu_2, T_2)$ are *isomorphic* if there is an invertible measure-preserving map ϕ between

them which takes orbits of T_1 to orbits of T_2 ‘almost everywhere’—that is, except perhaps for a set of points of measure zero (Ornstein and Weiss, 1991, 15–16). The intuitive idea is that the systems are isomorphic if each orbit of T_1 has a corresponding orbit of T_2 with identical probabilistic features, and the map ϕ tells us how to translate between them. To make the “almost everywhere” requirement explicit, we consider two subsets of M_1 and M_2 , \hat{M}_1 and \hat{M}_2 respectively, which differ from the full sets by a set of measure zero, and then demand that ϕ take orbits of T_1 to orbits of T_2 *everywhere* on \hat{M}_1 and \hat{M}_2 . The systems are called *manifestly isomorphic* if identical subsets \hat{M}_1 and \hat{M}_2 can be found.

If \hat{M}_1 and \hat{M}_2 are identical, then the two systems inhabit the same phase space, which means that they have the same possible outcomes. Since the systems are isomorphic, each trajectory in one system will have an analogue in the other, and corresponding bundles of trajectories will have the same probabilities. This sounds like the intuitive position, outlined above, that two systems could be observationally equivalent if they assigned the same probabilities to the same outcomes.

Along these lines, Werndl proposes the following definition of observational equivalence between deterministic and stochastic models. Suppose we want to determine whether a deterministic system (M, Σ_M, μ, T) and a stochastic process $\{Z_t; t \in \mathbb{Z}\}$ are observationally equivalent. Recall our assumption that we can only view the deterministic system through the finite-valued observation function Φ , which coarse-grains the phase space into a finite number of observable values. If our deterministic system is of the right type, specifically if it is totally ergodic, then applying a finite-valued observa-

tion function produces a non-trivial stochastic process (Werndl, 2009a, 235).³ In effect, when we coarse-grain the phase space with Φ we can no longer predict the system's next state with certainty, and this yields the non-trivial stochastic process $\{\Phi(T^t); t \in \mathbb{Z}\}$. Since manifest isomorphism is defined for deterministic systems, we consider the deterministic representation of these two stochastic processes. If the deterministic representation of the *derived* stochastic process $\{\Phi(T^t); t \in \mathbb{Z}\}$ is manifestly isomorphic to the deterministic representation of the *original* stochastic process $\{Z_t; t \in \mathbb{Z}\}$, then they have the same set of possible outcomes, and all trajectories in one have probabilistically equivalent analogues in the other. According to the proposal being entertained, this is exactly what it means for the process $\{Z_t; t \in \mathbb{Z}\}$ and the system (M, Σ_M, μ, T) (observed with Φ) to be observationally equivalent.

Formally, Werndl defines manifest isomorphic observational equivalence of a deterministic system (M, Σ_M, μ, T) and a stochastic process $\{Z_t; t \in \mathbb{Z}\}$ as follows:

The stationary stochastic process $\{Z_t; t \in \mathbb{Z}\}$ and the measure-preserving deterministic system (M, Σ_M, μ, T) , observed with Φ , are *observationally equivalent* if and only if the deterministic representation of $\{\Phi(T^t); t \in \mathbb{Z}\}$ is manifestly isomorphic to the deterministic representation of $\{Z_t; t \in \mathbb{Z}\}$. (Werndl, 2009a, 236)

³A measure-preserving transformation T is totally ergodic if T^n is ergodic for all integers $n > 0$ (Silva, 2008, 101). For more on ergodicity and its connection to randomness and chaos in deterministic systems, see Berkovitz et al. (2006) and Werndl (2009b)).

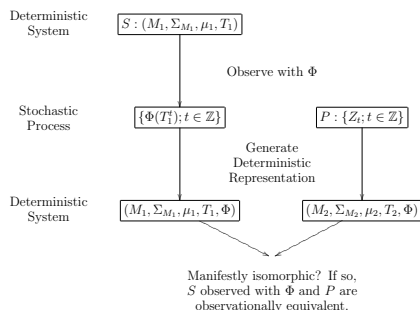


Figure 2: Schematic representation of Werndl’s proposed method of determining observational equivalence of deterministic and stochastic models.

A minor problem with manifest isomorphism is that the “if and only if” wording in Werndl’s definition rules out many model pairs that yield arbitrarily similar probabilistic predictions, and so might actually be observationally equivalent. Consider a deterministic Bernoulli system S_1 and a stochastic Bernoulli process P_1 with the same two outcomes, A and B . Let S_1 be shorthand for the system (M, Σ_M, μ, T) where M is the set of all bi-infinite sequences of A and B , let Σ_M be the σ -algebra generated by the appropriate cylinder sets on M , and let μ be the appropriate probability measure on Σ_M . Similarly, let P_1 be shorthand for $\{Z_t; t \in \mathbb{Z}\}$. Let the probabilities of the outcomes (A, B) in any individual trial be as follows:

$$S_1 : (0.5, 0.5) \tag{3}$$

$$P_1 : (0.50001, 0.49999) \tag{4}$$

Although not identical, intuitively, in many experimental situations the evolutions of S_1 and P_1 would seem to be observationally equivalent. However, since S_1 and P_1 have different KS-entropies, they cannot be manifestly isomorphic, and thus cannot be observationally equivalent according to the proposal under consideration. All this shows, of course, is that manifest

isomorphism will not do as a *necessary* condition for observational equivalence between these sorts of models. Since it might well be the case that it was never intended as such, we can avoid this problem by taking manifest isomorphism as a sufficient condition.

Yet, this modified proposal is also not sufficient for observational equivalence. Consider S_2 and P_2 , which are again two-outcome deterministic and stochastic Bernoulli models respectively, but this time set their probabilities as follows:

$$S_2 : (0.1, 0.9) \tag{5}$$

$$P_2 : (0.9, 0.1) \tag{6}$$

In order to satisfy the requirements for the proposed definition of observational equivalence, we need a stationary stochastic process $\{Z_t; t \in \mathbb{Z}\}$; a measure-preserving deterministic system (M, Σ_M, μ, T) observed with Φ ; and the deterministic representation of $\{\Phi(T^t); t \in \mathbb{Z}\}$ to be manifestly isomorphic to the deterministic representation of $\{Z_t; t \in \mathbb{Z}\}$.

The first condition is met by $P_2 = \{Z_t; t \in \mathbb{Z}\}$. The second condition is met by S_2 observed with Φ , since Werndl's proposition guarantees that this gives a stationary stochastic process $\{\Phi(T^t); t \in \mathbb{Z}\}$. For the third condition we need to show that the deterministic representations of $\{\Phi(T^t); t \in \mathbb{Z}\}$ and $\{Z_t; t \in \mathbb{Z}\}$ are manifestly isomorphic. By construction these models are isomorphic since they are Bernoulli and have the same KS-entropy. Their phase spaces also both consist of the set of all bi-infinite sequences of A and B . Since they are isomorphic and share the same phase space, these two deterministic representations are manifestly isomorphic and therefore observationally equivalent according to the proposal under consideration.

Yet, also by construction, S_2 and P_2 are intuitively *not* observationally equivalent since they assign very different probabilities to the same sequences of outcomes. S_2 and P_2 therefore stand in different relations to any set of evidence. In fact, by altering the probabilities of S_2 and P_2 , we can devise manifestly isomorphic models which assign *arbitrarily* different probabilities to any given observation. If observational equivalence for probabilistic models is supposed to be something like assigning the same (or similar) probabilities to the same outcome sequences, then manifest isomorphism is not well suited to the task.

A likely response to this counterexample would be to add the condition that two systems are observationally equivalent if the isomorphism map is the identity function $\phi(x) = x$.⁴ Since ϕ takes orbits of one system to orbits of the other almost everywhere, if ϕ is the identity function then almost all trajectories of the two deterministic representations are *identical*—and identity should certainly be sufficient for indistinguishability.

However, even if we do accept this modified definition, the inference from manifest isomorphism to observational equivalence requires non-mathematical assumptions, and so manifest isomorphism cannot be a *purely* mathematical definition of observational equivalence. For example, whether a deterministic model S and stochastic model P are manifestly isomorphic will depend in general on the specific finite-valued observation function Φ which is applied to S . But there are no mathematical facts which *compel* the choice of a particular Φ , and the selection of an appropriate Φ will be based on contextual

⁴Thanks to (Name redacted for blind review) for pointing this out to me.

and defeasible factors such as our confidence in our data, or physical theories about our measurement apparatus and the system under consideration. If we choose a strange or inappropriate Φ , such as the function Φ_6 which takes all elements of S 's phase space to the value 6, then S will be manifestly isomorphic to systems it is certainly not observationally equivalent to. Of course Φ_6 will be excluded from serious consideration in most circumstances, but this rejection is not a mathematical necessity. Rather, the decision to reject Φ_6 will be based on physical theories and beliefs about the system S , such as our expectation that it will deliver values other than 6. What this demonstrates is that the question of whether a manifestly isomorphic pair S and P are truly observationally equivalent depends on whether an appropriate Φ has been applied, but the appropriateness of a given Φ will vary depending on the circumstances, and there may not be a single correct—or at least undisputed—way of resolving this matter. While the manifest isomorphism of two models may be provable, the further inference to their observational equivalence will be based on contextual and defeasible factors.

To summarize, a purely mathematical definition of observational equivalence based on manifest isomorphism is in trouble if we accept the intuition that observational equivalence for probabilistic models should be something like assigning the same or similar probabilities to the same outcomes. Since manifestly isomorphic models can have radically different probability distributions, this definition picks out many wrong systems, and since non-manifestly isomorphic models can have arbitrarily similar probability distributions, it excludes many right ones. These problems can be avoided by restricting the claim to sufficiency and adding the requirement that $\phi(x) = x$.

However, while this modified manifest isomorphism is a provable relation, the inference to observational equivalence will only be as reliable as the non-mathematical assumptions supporting it.

2.2. ε -Congruence

In 1991 the mathematicians Ornstein and Weiss introduced ε -congruence, which was presented as a well-defined notion of observational equivalence (Ornstein and Weiss, 1991, 23). ε -congruence has been an influential concept, and some have argued that, since it entails observational equivalence, it has important implications for the metaphysical thesis of determinism (Suppes, 1993; Suppes and de Barros, 1996). In this section I will begin by outlining ε -congruence and the arguments purporting to establish it as observational equivalence. I will then argue that this view is untenable, since two dynamical systems can be ε -congruent yet observationally distinguishable. ε -congruence plus some extra conditions may be more feasible as a definition of observational equivalence, but, as with manifest isomorphism, these additional conditions will generally be fallible physical hypotheses, and the inference to observational equivalence will no longer be deductively certain.

Take two deterministic measure-preserving dynamical systems, associated with transformations T_1 and T_2 respectively, that act on the same phase space M . We introduce a *metric*, which is a function that gives the distance between any two points in M . Recall that T_1 and T_2 are isomorphic if there is a measure-preserving map ϕ which takes orbits of T_1 to orbits of T_2 almost everywhere. ε -congruence requires that two systems be isomorphic but also puts restrictions on their geometrical and statistical properties. Ornstein and Weiss give the following definition of ε -congruence for two systems inhabiting

the same phase space M :

We say that two measure-preserving [transformations] ... on the same compact metric space M are $[\varepsilon]$ -congruent if they are isomorphic and the map ϕ from M to M that implements the isomorphism moves the points in M by $< [\varepsilon]$ except for a set of points in M of measure $< [\varepsilon]$.⁵ (Ornstein and Weiss, 1991, 22–3)

More generally, let f_t and \bar{f}_t be flows on abstract measure spaces X and \bar{X} , and g and \bar{g} be functions from X and \bar{X} respectively to a metric space. Then we say that (f_t, X, g) and $(\bar{f}_t, \bar{X}, \bar{g})$ are ε -congruent if there is an invertible measure-preserving function ϕ such that $\phi f_t(x) = \bar{f}_t \phi(x)$ almost everywhere, and, letting d denote distance in the metric space, we have $d(g(x), \bar{g}(\phi x)) < \varepsilon$ everywhere except for in a set of measure less than ε (Ornstein and Weiss, 1991, 23).

Given the time-average interpretation of probability, the second part of this definition stipulates, roughly, that corresponding orbits are allowed to have at most distance ε between them ‘*most* of the time,’ but ‘ ε of the time’ they are allowed to be farther apart. ε is a parameter that ranges between zero and one, and the smaller ε is the more alike the two flows must be. If we set ε small but not too small, then we get an interesting notion of close-but-not-too-close: the corresponding trajectories of the two system are within some small distance ε of each other most of the time, but a proportionately small ε of the time they are allowed to differ by more.

Ornstein and Weiss thought this looked suggestive. They wrote:

⁵Notation changed for consistency.

If we agree that we cannot distinguish points in M that have distance $< [\varepsilon]$, and if we are willing to ignore events of probability less than $[\varepsilon]$ (experimental error), then $[\varepsilon]$ -congruent flows are indistinguishable. (Ornstein and Weiss, 1991, 23)

ε -congruence does seem promising as a mathematical definition of observational equivalence. For example, two Bernoulli systems with radically different probability distributions—such as those that gave rise to the counterexamples in section 2.1—could scarcely be ε -congruent for any small value of ε , since for the most part their trajectories will not be close to each other.

ε -congruence has attracted philosophical attention due to a theorem, proved by Ornstein and Weiss, which establishes that any deterministic Bernoulli system is, for all $\varepsilon > 0$, ε -congruent to some (generally ε -dependent) stochastic process (Ornstein and Weiss, 1991, 39). If to be ε -congruent is to be observationally equivalent, and if we accept Ornstein and Weiss's interpretation of ε -congruence, then this theorem entails that there are classical deterministic models which are indistinguishable from stochastic models *at all observation levels*.⁶

The ε -congruence theorem concerns deterministic Bernoulli systems and stochastic semi-Markov processes. A Bernoulli system, recall, is a deterministic measure-preserving system whose behaviour is chaotic and extremely unpredictable. A semi-Markov process is a stochastic process which remains in one of a finite number of states for a period of time and then jumps to

⁶Ornstein and Weiss's result has been philosophically influential, but other authors have argued for similar conclusions (e.g. Werndl (2011)).

a new state, where both the time between jumps and the transition probabilities are state dependent. Since the process is stochastic, the result of this jump cannot be predicted with certainty. Ornstein and Weiss's theorem proves that every deterministic Bernoulli system on a manifold M is, for every $\varepsilon > 0$, ε -congruent to some stochastic semi-Markov process on M , where this stochastic process will in general depend on the value of ε (Ornstein and Weiss, 1991, 39).

We might expect deterministic and stochastic models to behave very differently, yet the ε -congruence result appears to imply arbitrarily similar behaviour in some cases. This is surprising, to say the least, and Ornstein and Weiss suggested the following interpretation:

This may mean that there is no philosophical distinction between processes governed by roulette wheels and processes governed by Newton's laws ... we are comparing, in a strong sense, Newton's laws and coin flipping. (Ornstein and Weiss, 1991, 39)

Ornstein and Weiss did not develop this idea any further, but several philosophers have explored its implications. Patrick Suppes accepted that ε -congruence implies observational equivalence, and argued that this leads to a strong form of underdetermination wherein any thesis concerning the deterministic or indeterministic nature of the world must necessarily "transcend experience" (Suppes, 1993; Suppes and de Barros, 1996). In a reply to Suppes, John Winnie conceded that while the ε -congruence results show that deterministic Bernoulli and stochastic Markov models are observationally equivalent, there are *inductive* reasons for preferring the deterministic model, since it "outstrips any single Markov model in its conceptual and predictive power"

(Winnie, 1998, 317). Entering this debate here would take us far afield, but both authors accept as a starting premise that ε -congruence implies observational equivalence.

It is surprising, given the potentially far-reaching ramifications of Suppes and Winnie's debate, that very little attention has been given to the question of whether ε -congruence is in fact a sufficient condition for observational equivalence. In the remainder of this section I will consider the two main justifications of this claim in the literature, the first from Ornstein and Weiss's original paper, and a more recent and thorough treatment by Charlotte Werndl. I will argue that these justifications are problematic, and that ε -congruence, like manifest isomorphism, is susceptible to counterexamples which show that it cannot be sufficient for observational equivalence. I will close by suggesting that ε -congruence can be made more adequate by imposing additional conditions; however, as with manifest isomorphism, these conditions will in general be based on defeasible physical hypotheses.

To begin, Ornstein and Weiss's interpretation of ε -congruence risks inappropriately conflating an experiment's precision with its accuracy. Their interpretation, recall, is that if we cannot distinguish measurements within ε of each other, and if we ignore events of probability less than ε , then ε -congruence is observational equivalence. However, using the same parameter ε to quantify both precision and accuracy is difficult to motivate since we expect these factors to vary independently, and sometimes to differ quite greatly. If we use only one variable to account for both types of inexactitude, we seem committed to accepting greater uncertainty in each measurement if we eliminate more measurements as outliers. This will often be unwarranted.

If we use a precision instrument in a noisy environment, each measurement may be quite exact even if we choose to disregard a great many measurements as due to external noise factors. A delicate acoustical experiment, for example, will detect a great deal of outliers if it is located in a bowling alley; but even if we ignore a large proportion ε of its results, it would be unwarranted to consider each individual measurement correspondingly inexact. In cases where precision and accuracy differ, Ornstein and Weiss's interpretation of ε -congruence seems not to apply.

Werndl's explication of ε -congruence overcomes this difficulty by making ε dependent on two other quantities. First, she says, let ε_1 be the minimum distance at which states of the deterministic system can be distinguished. Then, note that "in practice, for sufficiently small ε_2 , one will not be able to observe differences in probabilities of less than ε_2 " (Werndl, 2009a, 238). Presumably ε_1 and ε_2 will be determined on an experiment-by-experiment basis, depending on the situation at hand. Now let ε be smaller than ε_1 and ε_2 . Then, claims Werndl, two models "give the same predictions at observation level ε " if their solutions can be put into one-to-one correspondence in such a way that at each time point they are less than ε apart, except for a set whose probability is smaller than ε . In other words, ε -congruence is indistinguishability at observation level ε .

Werndl's finer-grained approach is an improvement, but problems still arise when we try to set ε based on ε_1 and ε_2 . If we set ε greater than both ε_1 and ε_2 , the ε -congruence requirement will be less restrictive than the most restrictive condition imposed by the details of the experiment. This can result in systems being ε -congruent, and thus mislabelled as observationally

equivalent, even if they are quite clearly discriminable. Werndl prudently guards against this, and advises us instead to choose ε smaller than both ε_1 and ε_2 . Any variations between the two models must then occur below both thresholds of detectability. However, ε is more restrictive than the least restrictive condition imposed by the actual experiment, and consequently this condition will fail to identify some intuitively indistinguishable systems as observationally equivalent. Werndl’s conservative choice of ε will avoid misclassifying distinguishable systems as observationally equivalent, but only at the cost of denying ε -congruence the status of a necessary condition for observational equivalence.

It might be objected that picking on ε ’s relation to ε_1 and ε_2 is unfair, since arguably these finer-grained quantities underlie the mathematical discussion, but using a coarser-grained ε makes the proofs easier. Furthermore, in the *specific* case of the Orenstein-Weiss theorem, setting $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$ may actually be an acceptable heuristic.⁷ This is fine and well, but the present concern is interpreting this move in the light of a general notion of observational equivalence. If ε -congruence based on Werndl’s finer-grained ε_1 and ε_2 is not necessary for observational equivalence, then in order to be generally applicable it must at least be sufficient. Unfortunately, as I will now argue, this is not the case.

ε -congruence cannot be a sufficient condition for observational equivalence because the set of points where two systems differ by more than ε (hereafter called the ε -set, for brevity’s sake) is restricted only in its mea-

⁷Thanks to (Name redacted for blind review) for stressing this point.

sure, and not in its distribution. This means that two ε -congruent models can have ε -sets that differ in empirically meaningful ways. If the trajectories of two systems always remain within ε of each other they will be indistinguishable, and so observationally equivalent. If two systems do have an ε -set, then their trajectories differ at some points by more than ε , but we can still call them observationally equivalent if we write these variations off as random error. But if the trajectories of two systems differ at *extremely regular* intervals by *extremely large* amounts, then it is hard to see how they could be observationally equivalent. Two audio recordings of gently hissing white noise may reasonably be regarded as indistinguishable even if they are not identical. However, if one recording also includes the distinct and regular ticking of a clock while the other does not, any claim of their observational equivalence becomes very suspect.

This intuitive argument can be made more precise, and I will now construct two systems S_1 and S_2 that are distinguishable despite meeting all the technical requirements for ε -congruence. Let M be the section of the Cartesian plane $(0, 1] \times (0, 1]$ with opposite edges identified, let Σ_M be all Lebesgue-measurable subsets of M , and let μ be the Lebesgue measure. For simplicity, S_1 is a very boring system whose trajectories are straight lines moving constantly to the right. The trajectories of S_2 also move constantly rightward, but are permitted to deviate smoothly from straight lines. Let all the trajectories of S_2 be the same but shifted up or down on the y -axis, so the phase space of S_2 is filled with a stack of similar trajectories. Let S_1 be the system (M, Σ_M, μ, T_1) with horizontal trajectories given by the following




Figure 3: The bump function $x = e^{-\frac{1}{1-x^2}}$.

transformation T_1 :

$$T_1(x, y) = ((x + \tau) \bmod 1, y) \quad (7)$$

where τ is a parameter that determines the ‘speed’ with which trajectories of S_1 move across the phase space. Let $P(x)$ be a perturbing function. Many options present themselves, but here we will use a *bump function*, a continuous curve with continuous derivatives of all orders. A standard bump function is defined as follows:

$$B(x) = \begin{cases} e^{-\frac{1}{1-x^2}} & \text{if } |x| < 1 \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

In appearance $B(x)$ is reminiscent of a Gaussian curve, except it smoothly approaches and meets the line $y = 0$ at $x = \pm 1$ (see fig. 3). Here we will use the custom bump function $P(x)$. Since ε_1 and ε_2 can be quite different, to generate a counterexample we need two systems whose ε -set has measure smaller than the minimum of ε_1 and ε_2 , but whose largest deviations are greater than the detectability threshold ε_1 . Following Werndl, we choose a conservative value for ε , and set $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$. $P(x)$ is then defined as




Figure 4: Graph of $P(x)$ when $\varepsilon = 0.4$ and $\varepsilon_1 = 0.5$. $P(x)$ is below the dotted line $y = 0.4$ most of the time, but climbs above the dashed threshold of detectability $y = 0.5$ on an interval of width strictly less than 0.4.

follows:

$$P(x) = \begin{cases} \varepsilon(1 - \varepsilon) + \varepsilon_1 e^{\left(1 - \frac{\varepsilon^2}{\varepsilon^2 - 4(x-0.5)^2}\right)} & \text{if } |x - 0.5| < \frac{\varepsilon}{2}, \\ \varepsilon(1 - \varepsilon) & \text{otherwise.} \end{cases} \quad (9)$$

$P(x)$ provides a ‘bump’ of width ε and height $\varepsilon(1 - \varepsilon) + \varepsilon_1$ centred around $x = 0.5$. Outside this interval the curve connects smoothly to a horizontal line of height $\varepsilon(1 - \varepsilon)$. We can generate a system S_2 , given by the quadruple (M, Σ_M, μ, T_2) , by perturbing the transformation operator of S_1 with $P(x)$ as follows:

$$T_2(x, y) = ((x + \tau) \bmod 1, (y - P(x) + P((x + \tau) \bmod 1)) \bmod 1) \quad (10)$$

Trajectories of S_1 are mapped to trajectories of S_2 by the following ϕ :

$$\phi(x, y) = (x, (y + P(x)) \bmod 1) \quad (11)$$

Conversely, ϕ^{-1} takes trajectories of S_2 back to S_1 :

$$\phi^{-1}(x, y) = (x, (y - P(x)) \bmod 1) \quad (12)$$

The determinants of the Jacobians of $\phi(x, y)$ and its inverse are 1 and do not depend on x , y , ε_1 , or ε_2 , so $\phi(x, y)$ and its inverse are measure-preserving at all points and for all values of ε_1 and ε_2 . S_1 and S_2 are therefore isomorphic.

To establish ε -congruence, we need to show that $\phi(x, y)$ moves trajectories by less than ε , except for in a region of measure less than ε , where $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$. Since the situation will be the same for all trajectories modulo a vertical translation in the phase space, consider the trajectory of S_1 that moves uniformly across the phase space along the line $y = 0$. The corresponding trajectory of S_2 spends most of its time moving along the horizontal line $y = \varepsilon(1 - \varepsilon)$, and so lies within ε of S_1 . Within the bump of width ε , the trajectory increases to a height of $\varepsilon(1 - \varepsilon) + \varepsilon_1$, which is greater than ε_1 , and so represents a detectable deviation from the trajectory of S_1 . Since the bump began slightly below $y = \varepsilon$, the portion of the bump above this line, which is this trajectory's contribution to the ε -set, will be of width strictly less than ε . Similar considerations apply for all trajectories, and so the total ε -set has measure less than ε . Therefore, no matter which of ε_1 and ε_2 is smaller, S_1 and S_2 are isomorphic, inhabit the same phase space, and the measure-preserving map $\phi(x, y)$ that takes trajectories of S_1 to trajectories of S_2 moves orbits by less than ε except for a set of measure less than ε . S_1 and S_2 are therefore ε -congruent, but have a regular and detectable difference.

The function $P(x)$ is not terribly exciting as it stands, but it can be modified to generate a more pathological perturbation $P^*(x)$ (see fig 5). First, the size of the bump can be made arbitrarily large by increasing the coefficient of the exponential term. ε_1 was chosen here because it ensures




Figure 5: Graph of $P^*(x)$ with five peaks, again with $\varepsilon = 0.4$ (dotted line) and $\varepsilon_1 = 0.5$ (dashed line).

detectability and seems to fit with the general spirit of the proposal, but the systems will remain ε -congruent under a perturbation by an arbitrarily large bump. Second, while $P(x)$ consists of only one bump, the number of bumps in $P^*(x)$ can be made an arbitrarily large number n by compressing the original function $P(x)$ to a width of $1/n$ and repeating it n times. Since the width of the observable portion of the original bump is less than ε , the width of the observable portions of each shrunken bump will be less than ε/n , and the combined width of all n observable bumps will still be less than ε . If we perturb S_1 with such a $P^*(x)$, the result will be an ε -congruent S_2 with arbitrarily many arbitrarily large deviations.⁸

Therefore no matter which of ε_1 or ε_2 is smaller, two systems can be ε -congruent and yet differ in empirically meaningful, observationally detectable ways. Since in both cases the trajectories of S_2 are those of S_1 perturbed

⁸Although not proved here, this result should generalize to an arbitrary S_1 , further problematizing any contextless interpretation of the ε -congruence of two mathematical models.

by a regular function, the trajectories of S_2 will vary measurably from those of S_1 at regular intervals. Since both the number and size of the bumps in the perturbing function $P^*(x)$ can be made arbitrarily large, the number and size of the of detectable variations between S_1 and S_2 can be made arbitrarily large. It is highly counter intuitive to say that, in the absence of any other considerations, two mathematical models could be observationally equivalent when they diverge detectably and systematically. Some contextual or theoretical explanation would need to be given for why these observations differ systematically yet still count as observationally equivalent, but any explanation will necessarily go beyond the mathematics. The conclusion is that ε -congruence alone cannot be sufficient for observational equivalence, since it will, at least in some cases, need to be supplemented with a physical theory, model, or hypothesis about the systems in question.

Perhaps for a sufficiently small ε we can expect not to observe any members of the ε -set, and so it can be safely ignored. However, this cannot be guaranteed. In many cases, and particularly in particle physics, the number of measurements performed may be quite enormous, making it highly probable that a member of the ε -set will be observed. Furthermore, these improbable outlying events may be important. This is, in fact, often the point, and many experiments are designed specifically to detect such low-frequency events.

Perhaps ε -congruence could be strengthened with the requirement that the ε -set be distributed randomly in some sense, to match our expectations about experimental noise. Indeed, since the systems considered by Ornstein and Weiss are Bernoulli, and therefore strongly chaotic, it already seems

implausible, although not impossible, that the outliers in the ε -set will be distributed irregularly. This may be feasible, but there are two possible problems. First, an explicit randomness requirement on the ε -set would complicate the mathematics, and there is no guarantee that any of the interesting results about deterministic and stochastic processes—which, recall, are what sparked philosophical interest in the subject in the first place—would still hold with this modified definition. Second, there are many different definitions of randomness, and the justification of any choice would likely have to take into account the theoretical characteristics of the system and the experimental context at hand. This would introduce a large amount of context sensitivity into the definition, once again eliminating its purely mathematical and general nature.

Given these considerations, the prospects for using ε -congruence as a *purely* mathematical definition of observational equivalence seem dim. Strict ε -congruence permits of counterexamples, and resolving these counterexamples will introduce non-mathematical considerations. The mere ε -congruence of two models therefore does not entail their observational equivalence. For ε -congruence to be an adequate criterion of observational equivalence, it must be supplemented by hypotheses about the physical systems under study, the assumptions embedded in our models, and the kinds of data discrepancies we are willing to attribute to chance error events.

3. Conclusion

Deterministic and stochastic models can be manifestly isomorphic and ε -congruent, but these relations alone are not adequate formalizations of ob-

servational equivalence since the strict mathematical requirements of each can be met by distinguishable systems. Nor can either relation easily be made into purely mathematical sufficient conditions, since natural attempts to strengthen them will introduce non-mathematical physical hypotheses. Thus, it seems that whether two models are observationally equivalent or not will be a context-sensitive judgement based on physical hypotheses, and neither manifest isomorphism nor ε -congruence can deal with these sorts of experimental vagaries in a strict, axiomatic way. Plausibly, *any* purely algorithmic approach to observational equivalence will risk ignoring contextual subtleties.

This only goes to show that Werndl's purely mathematical definitions are incomplete, not that they are wrong, and indeed I believe manifest isomorphism and ε -congruence can guide us towards a more nuanced understanding of observational equivalence between certain types of mathematical models. By making it clear which parts of our arguments for observational equivalence *are* provable, Werndl's definitions can likewise help us to determine which parts are not. Armed with this knowledge, anyone wishing to advance, or to dispute, an argument for observational equivalence on the basis of manifest isomorphism or ε -congruence can do so in a more productive way. If these arguments also contain non-mathematical components, then our judgements of observational equivalence will only be as strong as the theoretical assumptions underpinning them; but this seems appropriate for a scientific concept.

References

- Berkovitz, J., Frigg, R., Kronz, F., 2006. 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.
- Doob, J. L., 1953. *Stochastic processes*. Wiley, New York.
- Judd, K., Smith, L., 2004. Indistinguishable states ii. Imperfect model scenarios. *Physica D Nonlinear Phenomena* 196 (3-4), 224–242.
- Ornstein, D., 1970. Bernoulli shifts with the same entropy are isomorphic. *Advances in Mathematics* 4 (3), 337–352.
- Ornstein, D., Weiss, B., 1991. Statistical properties of chaotic systems. *Bulletin of the American Mathematical Society* 24 (1), 11–116.
- Silva, C. E., 2008. *Invitation to Ergodic Theory*. American Mathematical Society, Providence, RI.
- Suppes, P., 1993. The transcendental character of determinism. *Midwest Studies in Philosophy* 18, 242–57.
- Suppes, P., de Barros, J., 1996. Photons, billiards, and chaos. In: Weingartner, P., Schurz, G. (Eds.), *Law and Prediction in the Light of Chaos Research*. Springer Berlin / Heidelberg, Berlin, pp. 189–201.
- Werndl, C., 2009a. Are deterministic descriptions and indeterministic descriptions observationally equivalent? *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics* 40 (3), 232–242.

- Werndl, C., 2009b. What are the new implications of chaos for unpredictability? *The British Journal for the Philosophy of Science* 60 (1), 195–220.
- Werndl, C., 2011. On the observational equivalence of continuous-time deterministic and indeterministic descriptions. *European Journal for Philosophy of Science* 3 (1), 193–225.
- Winnie, J. A., 1998. *The Cosmos of Science: Essays of Exploration*. University of Pittsburgh Press, Pittsburgh, Pa., Ch. Deterministic chaos and the nature of chance, pp. 299–324.