Informational and Causal Architecture of
Discrete-Time Renewal Processes

Sarah Marzen[1, *] and James P. Crutchfield[2, †]

[1]*Redwood Center for Theoretical Neuroscience and Department of Physics,*
University of California at Berkeley, Berkeley, CA 94720-5800
[2]*Complexity Sciences Center and Department of Physics,*
University of California at Davis, One Shields Avenue, Davis, CA 95616
(Dated: July 1, 2015)

Renewal processes are broadly used to model stochastic behavior consisting of isolated events separated by periods of quiescence, whose durations are specified by a given probability law. Here, we identify the minimal sufficient statistic for their prediction (the set of causal states), calculate the historical memory capacity required to store those states (statistical complexity), delineate what information is predictable (excess entropy), and decompose the entropy of a single measurement into that shared with the past, future, or both. The causal state equivalence relation defines a new subclass of renewal processes with a finite number of causal states despite having an unbounded interevent count distribution. We use these formulae to analyze the output of the parametrized Simple Nonunifilar Source, generated by a simple two-state hidden Markov model, but with an infinite-state ϵ-machine presentation. All in all, the results lay the groundwork for analyzing processes with infinite statistical complexity and infinite excess entropy.

I. INTRODUCTION

Stationary renewal processes are widely used, analytically tractable, compact models of an important class of point processes [1–4]. Realizations consist of sequences of events—e.g., neuronal spikes or earthquakes—separated by epochs of quiescence, the lengths of which are drawn independently from the same interevent distribution. Renewal processes on their own have a long history and, due to their offering a parsimonious mechanism, often are implicated in highly complex behavior [5–10]. Additionally, understanding more complicated processes [11–14] requires fully analyzing renewal processes and their generalizations.

As done here and elsewhere [15], analyzing them in-depth from a structural information viewpoint yields new statistical signatures of apparent high complexity—long-range statistical dependence, memory, and internal structure. To that end, we derive the causal-state minimal sufficient statistics—the ϵ-machine—for renewal processes and then derive new formulae for their various information measures in terms the interevent count distribution. The result is a thorough-going analysis of their *information architecture*—a shorthand referring to a collection of measures that together quantify key process properties:

predictability, difficulty of prediction, inherent randomness, memory, and Markovity, and the like. The measures include:

- the *statistical complexity* C_μ, which quantifies the historical memory that must be stored in order to predict a process's future;
- the *entropy rate* h_μ, which quantifies a process' inherent randomness as the uncertainty in the next observation even given that we can predict as well as possible;
- the *excess entropy* **E**, which quantifies how much of a process's future is predictable in terms of the mutual information between its past and future;
- the *bound information* b_μ, which identifies the portion of the inherent randomness (h_μ) that affects a process's future in terms of the information in the next observation shared with the future, above and beyond that of the entire past; and
- the *elusive information* σ_μ, which quantifies a process's deviation from Markovity as the mutual information between the past and future conditioned on the present.

Analyzing a process in this way gives a more detailed understanding of its structure and stochasticity. Beyond this, these information measures are key to finding limits to a process's optimal lossy predictive features [16–19], designing action policies for intelligent autonomous agents [20], and quantifying whether or not a given process has one or another kind of infinite memory [21–23].

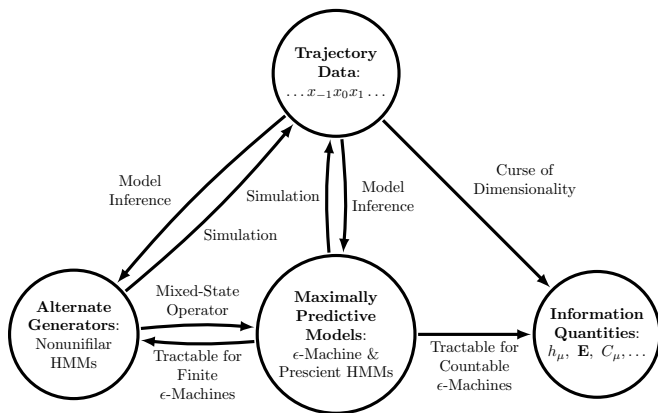* smarzen@berkeley.edu
† chaos@ucdavis.edu

FIG. 1. **The role of maximally predictive (prescient) models:** Estimating information measures directly from trajectory data encounters a curse of dimensionality or, in other words, severe undersampling. Instead, one can calculate information measures in closed-form from (inferred) maximally predictive models [24]. Alternate generative models that are *not* maximally predictive cannot be used directly, as Blackwell showed in the 1950s [25].

While it is certainly possible to numerically estimate information measures directly from trajectory data, statistical methods generally encounter a *curse of dimensionality* when a renewal process has long-range temporal correlations since the number of typical trajectories grows exponentially (at entropy rate h_μ). Alternatively, we gain substantial advantages by first building a maximally predictive model of a process (e.g., using Bayesian inference [26]) and then using that model to calculate information measures (e.g., using recently available closed-form expressions when the model is finite [24]). Mathematicians have known for over a half century [25] that alternative models that are not maximally predictive are inadequate for such calculations. Thus, maximally predictive models are critical. Figure 1 depicts the overall procedure just outlined, highlighting their important role. Here, extending the benefits of this procedure, we determine formulae for the information measures mentioned above and the appropriate model structures for a class of processes that require countably infinite models—the ubiquitous renewal processes.

Our development requires familiarity with computational mechanics [27]. Those disinterested in its methods, but who wish to use the results, can skip to Figs. 3(a)-3(d) and Table I. A pedagogical example is provided in Sec. V. Two sequels will use the results to examine the limit of infinitesimal time resolution for information in neural spike trains [28] and the conditions under which renewal processes have infinite excess entropy [29].

The development is organized as follows. Section II provides a quick introduction to computational mechanics and prediction-related information measures of stationary time series. Section III identifies the causal states (in both forward and reverse time), the statistical complexity, and the ϵ-machine of discrete-time stationary renewal processes. Section IV calculates the information architecture and predictable information of a discrete-time stationary renewal process. Section V calculates these information-theoretic measures explicitly for the parametrized Simple Nonunifilar Source, a simple two-state Hidden Markov Model with a countable infinity of causal states. Finally, Sec. VI summarizes the results and lessons, giving a view to future directions and mathematical and empirical challenges.

II. BACKGROUND

We first describe renewal processes, then introduce a small piece of information theory, review the definition of process structure, and finally recall several information-theoretic measures designed to capture organization in structured processes.

A. Renewal Processes

We are interested in a system's immanent, possibly emergent, properties. To this end we focus on behaviors and not, for example, particular equations of motion or particular forms of stochastic differential or difference equation. The latter are important in applications because they are generators of behavior, as we will see shortly in Sec. II C. As Fig. 1 explains, for a given process, some of its generators facilitate calculating key properties. Others lead to complicated calculations and others still cannot be used at all.

As a result, our main object of study is a *process* $\mathcal{P}$: the list of all of a system's behaviors or realizations $\{\ldots, x_{-2}, x_{-1}, x_0, x_1, \ldots\}$ as specified by their measure $\mu(\ldots, X_{-2}, X_{-1}, X_0, X_1, \ldots)$. We denote a contiguous chain of random variables as $X_{0:L} = X_0 X_1 \cdots X_{L-1}$. Left indices are inclusive; right, exclusive. We suppress indices that are infinite. In this setting, the *present* X_0 is the random variable measured at $t = 0$, the *past* is the chain $X_{:0} = \ldots X_{-2} X_{-1}$ leading up the present, and the *future* is the chain following the present $X_{1:} = X_1 X_2 \cdots$. The joint probabilities $\Pr(X_{0:N})$ of sequences are determined by the measure of the corresponding cylinder sets: $\Pr(X_{0:N} = x_0 x_1 \ldots x_{N-1}) = \mu(\ldots, x_0, x_1, \ldots, x_{N-1}, \ldots)$. Finally, we assume a process is ergodic and stationary—$\Pr(X_{0:L}) = \Pr(X_{t:L+t})$ for all $t \in \mathbb{Z}$—and the observation values x_t range over a finite alphabet: $x \in \mathcal{A}$. In short, we work with *hidden Markov processes* [30].

Discrete-time stationary *renewal processes* here have binary observation alphabets $\mathcal{A} = \{0, 1\}$. Observation of the binary symbol 1 is called an *event*. The event *count* is the number of 0's between successive 1s. Counts n are i.i.d. random variables drawn from an *interevent distribution* $F(n)$, $n \geq 0$. We restrict ourselves to *persistent* renewal processes, such that the probability distribution function is normalized: $\sum_{n=0}^{\infty} F(n) = 1$. This translates into the processes being ergodic and stationary. We also define the *survival function* by $w(n) = \sum_{n'=n}^{\infty} F(n')$, and the expected interevent count is given by $\mu = \sum_{n=0}^{\infty} n F(n)$. We assume also that $\mu < \infty$. It is straightforward to check that $\sum_{n=0}^{\infty} w(n) = \mu + 1$.

Note the dual use of μ. On the one hand, it denotes the measure over sequences and, since it determines probabilities, it appears in names for informational quantities. On the other, it is a commonplace in renewal process theory that denotes mean rates. Fortunately, context easily distinguishes the meaning through the very different uses.

B. Process Unpredictability

The information or uncertainty in a process is often defined as the Shannon entropy $H[X_0]$ of a single symbol X_0 [31]:

$$H[X_0] = -\sum_{x \in \mathcal{A}} \Pr(X_0 = x) \log_2 \Pr(X_0 = x) . \qquad (1)$$

However, since we are interested in general complex processes—those with arbitrary dependence structure—we employ the *block entropy* to monitor information in long sequences:

$$
\begin{aligned}
H(L) &= H[X_{0:L}] \\
&= -\sum_{w \in \mathcal{A}^L} \Pr(X_{0:L} = w) \log_2 \Pr(X_{0:L} = w) .
\end{aligned}
$$

To measure a process's asymptotic per-symbol uncertainty one then uses the Shannon entropy rate:

$$h_\mu = \lim_{L \to \infty} \frac{H(L)}{L} ,$$

when the limit exists. (Here and elsewhere, μ reminds us that information quantities depend on the process's measure μ over sequences.) h_μ quantifies the rate at which a stochastic process generates information. Using standard informational identities, one sees that the entropy rate is also given by the conditional entropy:

$$h_\mu = \lim_{L \to \infty} H[X_0 | X_{-L:0}] . \qquad (2)$$

This form makes transparent its interpretation as the residual uncertainty in a measurement given the infinite past. As such, it is often employed as a measure of a process's degree of unpredictability.

C. Maximally Predictive Models

Forward-time causal states $\boldsymbol{\mathcal{S}}^+$ are minimal sufficient statistics for predicting a process's future [32, 33]. This follows from their definition—a *causal state* $\sigma^+ \in \boldsymbol{\mathcal{S}}^+$ is a sets of pasts grouped by the equivalence relation $\sim^+$:

$$
\begin{aligned}
x_{:0} &\sim^+ x'_{:0} \\
&\Leftrightarrow \Pr(X_{0:}|X_{:0} = x_{:0}) = \Pr(X_{0:}|X_{:0} = x'_{:0}) . \quad (3)
\end{aligned}
$$

So, $\boldsymbol{\mathcal{S}}^+$ is a set of classes—a coarse-graining of the uncountably infinite set of all pasts. At time t, we have the random variable $\mathcal{S}_t^+$ that takes values $\sigma^+ \in \boldsymbol{\mathcal{S}}^+$ and describes the *causal-state process* $\ldots, \mathcal{S}_{-1}^+, \mathcal{S}_0^+, \mathcal{S}_1^+, \ldots$. $\mathcal{S}_t^+$ is a partition of pasts $X_{:t}$ that, according to the indexing convention, does not include the present observation X_t. In addition to the set of pasts leading to it, a causal state σ_t^+ has an associated *future morph*—the conditional measure $\mu(X_{t:}|\sigma_t^+)$ of futures that can be generated from it. Moreover, each state σ_t^+ inherits a probability $\pi(\sigma_t^+)$ from the process's measure over pasts $\mu(X_{:t})$. The forward-time *statistical complexity* is defined as the Shannon entropy of the probability distribution over forward-time causal states [32]:

$$C_\mu^+ = H[\mathcal{S}_0^+] . \qquad (4)$$

A generative model is constructed out of the causal states by endowing the causal-state process with transitions:

$$T_{\sigma\sigma'}^{(x)} = \Pr(\mathcal{S}_{t+1}^+ = \sigma', X_t = x | \mathcal{S}_t^+ = \sigma) ,$$

that give the probability of generating the next symbol x and ending in the next state σ', if starting in state σ. (Residing in a state and generating a symbol do not occur simultaneously. Since symbols are generated during transitions there is, in effect, a half time-step difference in the indexes of the random variables X_t and $\mathcal{S}_t^+$. We suppress notating this.) To summarize, a process's *forward-time ϵ-machine* is the tuple $\{\mathcal{A}, \boldsymbol{\mathcal{S}}^+, \{T^{(x)} : x \in \mathcal{A}\}\}$.

For a discrete-time, discrete-alphabet process, the ϵ-machine is its minimal unifilar Hidden Markov Model (HMM) [32, 33]. (For general background on HMMs see [34–36].) Note that the causal state set can be finite, countable, or uncountable; the latter two cases can occur even for processes generated by finite-state HMMs. *Minimality* can be defined by either the smallest number of

states or the smallest entropy over states [33]. *Unifilarity* is a constraint on the transition matrices $T^{(x)}$ such that the next state σ' is determined by knowing the current state σ and the next symbol x. That is, if the transition exists, then $\Pr(\mathcal{S}_{t+1}^+|X_t = x, \mathcal{S}_t^+ = \sigma)$ has support on a single causal state.

While the ϵ-machine is a process's minimal, maximally predictive model, there can be alternative HMMs that are as predictive, but are not minimal. We refer to the maximally predictive property by referring to the ϵ-machine and these alternatives as *prescient*. The state and transition structure of a prescient model allow one to immediately calculate the entropy rate h_μ, for example. More generally, any statistic that gives the same (optimal) level of predictability, we call a *prescient statistic*.

A similar equivalence relation can be applied to find minimal sufficient statistics for retrodiction [37]. Futures are grouped together if they have equivalent conditional probability distributions over pasts:

$$x_{0:} \sim^- x'_{0:}$$
$$\Leftrightarrow \Pr(X_{:0}|X_{0:} = x_{0:}) = \Pr(X_{:0}|X_{0:} = x'_{0:}) . \quad (5)$$

A cluster of futures—a *reverse-time causal state*—defined by $\sim^-$ is denoted $\sigma^- \in \mathcal{S}^-$. Again, each σ^- inherits a probability $\pi(\sigma^-)$ from the measure over futures $\mu(X_{0:})$. And, the *reverse-time statistical complexity* is the Shannon entropy of the probability distribution over reverse-time causal states:

$$C_\mu^- = H[\mathcal{S}_0^-] . \quad (6)$$

In general, the forward and reverse-time statistical complexities are not equal [37, 38]. That is, different amounts of information must be stored from the past (future) to predict (retrodict). Their difference $\Xi = C_\mu^+ - C_\mu^-$ is a process's *causal irreversibility* and it reflects this statistical asymmetry.

Since we work with stationary processes in the following the time origin is arbitrary and so we drop the time index t when it is unnecessary.

D. Information Measures for Processes

Shannon's various information quantities—entropy, conditional entropy, mutual information, and the like—when applied to time series are functions of the joint distributions $\Pr(X_{0:L})$. Importantly, they define an algebra of information measures for a given set of random variables [39]. Reference [40] used this to show that the past and future partition the single-measurement entropy $H(X_0)$ into several distinct measure-theoretic

atoms. These include the *ephemeral information*:

$$r_\mu = H[X_0|X_{:0}, X_{1:}] , \quad (7)$$

which measures the uncertainty of the present knowing the past and future; the *bound information*:

$$b_\mu = I[X_0; X_{1:}|X_{:0}] , \quad (8)$$

which is the mutual information shared between present and future conditioned on past; and the *enigmatic information*:

$$q_\mu = I[X_0; X_{:0}; X_{1:}] , \quad (9)$$

which is the three-way mutual information between past, present, and future. Multi-way mutual informations are sometimes referred to as *co-informations* [40, 41] and, compared to Shannon entropies and two-way mutual information, can have counterintuitive properties, such as being negative.

For a stationary time series, the bound information is also the shared information between present and past conditioned on the future:

$$b_\mu = I[X_0; X_{:0}|X_{1:}]. \quad (10)$$

One can also consider the amount of predictable information not captured by the present:

$$\sigma_\mu = I[X_{:0}; X_{1:}|X_0] . \quad (11)$$

This is called the *elusive information* [42]. It measures the amount of past-future correlation not contained in the present. It is nonzero if the process *necessarily* has hidden states and is therefore quite sensitive to how the state space is observed or coarse grained.

The maximum amount of information in the future predictable from the past (or vice versa) is the *excess entropy*:

$$\mathbf{E} = I[X_{:0}; X_{0:}] .$$

It is symmetric in time and a lower bound on the stored informations C_μ^+ and C_μ^-. It is directly given by the information atoms above:

$$\mathbf{E} = b_\mu + \sigma_\mu + q_\mu . \quad (12)$$

The process's Shannon entropy rate h_μ—recall the form of Eq. (2)—can also be written as a sum of atoms:

$$h_\mu = H[X_0|X_{:0}]$$
$$= r_\mu + b_\mu .$$

Thus, a portion of the information (h_μ) a process spontaneously generates is thrown away (r_μ) and a portion is actively stored (b_μ). Putting these observations together gives the information architecture of a single measurement (Eq. (1)):

$$H[X_0] = q_\mu + 2b_\mu + r_\mu \ . \tag{13}$$

These identities can be used to determine r_μ, q_μ, and $\mathbf{E}$ from $H[X_0]$, b_μ, and σ_μ, for example.

We have a particular interest in when C_μ and $\mathbf{E}$ are infinite and so will investigate finite-time variants of causal states and finite-time estimates of statistical complexity and $\mathbf{E}$. For example, the latter is given by:

$$\mathbf{E}(M, N) = I[X_{-M:0}; X_{0:N}] \ . \tag{14}$$

If $\mathbf{E}$ is finite, then $\mathbf{E} = \lim_{M,N\to\infty} \mathbf{E}(M, N)$. When $\mathbf{E}$ is infinite, then the way in which $\mathbf{E}(M, N)$ diverges is one measure of a process' complexity [21, 43, 44]. Analogous, finite past-future (M, N)-parametrized equivalence relations lead to finite-time forward and reverse causal states and statistical complexities $C_\mu^+(M, N)$ and $C_\mu^-(M, N)$.

III.   CAUSAL ARCHITECTURE OF RENEWAL PROCESSES

It will be helpful pedagogically to anchor our theory in the contrast between two different, but still simple, renewal processes. One is the familiar "memoryless" Poisson process with rate λ. Its HMM generator, a biased coin, is shown at the left of Fig. 2. It has an interevent count distribution $F(n) = (1 - \lambda)\lambda^n$; a distribution with unbounded support. However, we notice in Fig. 2 that it is a unifilar model with a minimal number of states. So, in fact, this one-state machine *is* the ϵ-machine of a Poisson process. The rate at which it generates information is given by the entropy rate: $h_\mu = H(\lambda)$ bits per output symbol. (Here, $H(p)$ is the binary entropy function.) It also has a vanishing statistical complexity $C_\mu^+ = 0$ and so stores no historical information.

The second example is the *Simple Nonunifilar Source* (SNS) [45]; its HMM generator is shown in Fig. 2(b). Transitions from state B are unifilar, but transitions from state A are not. In fact, a little reflection shows that the time series produced by the SNS is a discrete-time renewal process. Once we observe the "event" $x_t = 1$, we know the internal model state to be $\sigma_{t+1} = A$, so successive interevent counts are completely uncorrelated.

The SNS generator is not an ϵ-machine and, moreover, it cannot be used to calculate the process's information per output symbol (entropy rate). If we can only see 0's and 1's, we will usually be uncertain as to whether we
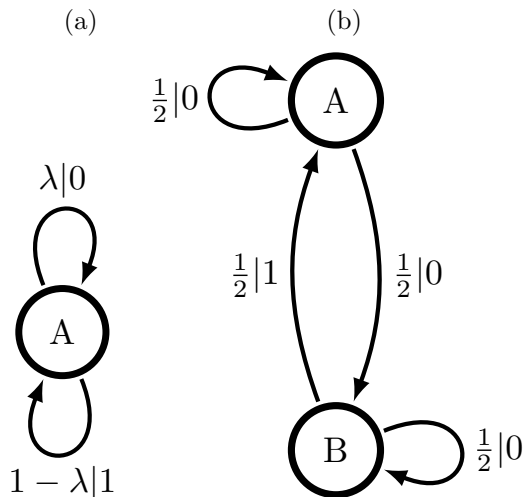

FIG. 2. (a) Minimal generative model for the Poisson process with rate λ. (b) A generator for the Simple Nonunifilar Source (SNS). Both generate a stationary renewal process. Transition labels $p|s$ denote probability p of taking a transition and emitting symbol s.

are in state A or state B, so this generative model is not maximally predictive. How can we calculate this basic quantity? And, if we cannot use the two-state generator, how many states are required and what is their transition dynamic? The following uses computational mechanics to answer these and a number of related questions. To aid readability, though, we sequester most all of the detailed calculations and proofs in App. A.

We start with a simple Lemma that follows directly from the definitions of a renewal process and the causal states. It allows us to introduce notation that simplifies the development.

Lemma 1. *The count since last event is a prescient statistic of a discrete-time stationary renewal process.*

That is, if we remember only the number of counts since the last event and nothing prior, we can predict the future as well as if we had memorized the entire past. Specifically, a *prescient state* $\mathcal{R}$ is a function of the past such that:

$$H[X_{0:}|X_{:0}] = H[X_{0:}|\mathcal{R}] \ .$$

Causal states can be written as unions of prescient states [33]. We start with a definition that helps to characterize the converse; i.e., when the prescient states of Lemma 1 are also causal states.

To ground our intuition, recall that Poisson processes are "memoryless". This may seem counterintuitive, if viewed from a parameter estimation point of view. After all, if observing longer pasts, one makes better and better estimates of the Poisson rate. However, finite data fluc-

tuations in estimating model parameters are irrelevant to the present mathematical setting *unless* the parameters are themselves random variables, as in Ref. [44]. This is not our setting here: the parameters are fixed. In fact, we restrict ourselves to studying ergodic processes, in which the conditional probability distributions of futures given pasts of a Poisson process are independent of the past.

We therefore expect the prescient states in Lemma 1 to fail to be causal states precisely when the interevent distribution is similar to that of a Poisson renewal process. This intuition is made precise by Def. 2.

Definition 1. *A* Δ-Poisson *process has an interevent distribution*

$$F(n) = F(n \bmod \Delta)\, \lambda^{\lfloor n/\Delta \rfloor}\,,$$

for all n and some $\lambda > 0$. If this statement holds for multiple $\Delta \geq 1$, then we choose the smallest possible Δ.

Definition 2. *A $(\tilde{n}, \Delta)$ eventually Δ-Poisson process has an interevent distribution that is Δ-Poisson for all $n \geq \tilde{n}$:*

$$F(\tilde{n} + k\Delta + m) = \lambda^k F(\tilde{n} + m)\,,$$

for all $0 \leq m < \Delta$, for all $k \geq 0$, and for some $\lambda > 0$. If this statement holds for multiple $\Delta \geq 1$ and multiple $\tilde{n}$, then we choose the smallest possible Δ and the smallest possible $\tilde{n}$.

If this statement holds for multiple $\Delta \geq 1$ and multiple n, *first* we choose the smallest possible Δ and *then* select the smallest possible n *for that* Δ. Order matters here.

Thus, a Poisson process is a Δ-Poisson process with $\Delta = 1$ and an eventually Δ-Poisson process with $\Delta = 1$ and $\tilde{n} = 0$. Moreover, we will now show that at some finite $\tilde{n}$, any renewal process is either (i) Poisson, if $\Delta = 1$, or (ii) a combination of several Poisson processes, if $\Delta > 1$.

Why identify new classes of renewal process? In short, renewal processes that are similar to, but not the same as, the Poisson process do not have an infinite number of causal states. The particular condition for when they do not is given by the eventually Δ-Poisson definition. Notably, this new class is what emerged, rather unexpectedly, by applying the causal-state equivalence relation $\sim^+$ to renewal processes. The resulting insight is that general renewal processes, after some number of counts (the "eventually" part) and after some coarse-graining of counts (the Δ part), behave like a Poisson process.

With these definitions in hand, we can proceed to identify the causal architecture of discrete-time stationary renewal processes.

Theorem 1. *(a) The forward-time causal states of a discrete-time stationary renewal process that is not even-* *tually Δ-Poisson are groupings of pasts with the same count since last event. (b) The forward-time causal states of a discrete-time eventually Δ-Poisson stationary renewal process are groupings of pasts with the same count since last event up until $\tilde{n}$ and pasts whose count n since last event are in the same equivalence class as $\tilde{n}$ modulo Δ.*

The Poisson process, as an eventually Δ-Poisson with $\tilde{n} = 0$ and $\Delta = 1$, is represented by the one-state ϵ-machine despite the unbounded support of its interevent count distribution. Unlike most processes, the Poisson process' ϵ-machine is the same as its generative model shown in Fig. 2(a).

The SNS, on the other hand, has an interevent count distribution that is not eventually Δ-Poisson. According to Thm. 1, then, the SNS has a countable infinity of causal states despite its simple two-state generative model in Fig. 2(b). Compare Fig. 3(a). Each causal state corresponds to a different probability distribution over the internal states A and B. These internal state distributions are the *mixed states* of Ref. [46]. Observing more 0's, one becomes increasingly convinced that the internal state is B. For maximal predictive power, however, we must track the probability that the process is still in state A. Both Fig. 3(a) and Fig. 2(b) are "minimally complex" models of the same process, but with different definitions of model complexity. We return to this point in Sec. V.

Appendix A makes the statements in Thm. 1 precisely. The main result is that causal states are sensitive to two features: (i) eventually Δ-Poisson structure in the interevent distribution and (ii) the boundedness of $F(n)$'s support. If the support is bounded, then there are a finite number of causal states rather than a countable infinity of causal states. Similarly, if $F(n)$ has Δ-Poisson tails, then there are a finite number of causal states despite the support of $F(n)$ having no bound. Nonetheless, one can say that the generic discrete-time stationary renewal process has a countable infinity of causal states.

Finding the probability distribution over these causal states is straightforwardly related to the survival-time distribution $w(n)$ and the mean interevent interval μ, since the probability of observing at least n counts since last event is $w(n)$. Hence, the probability of seeing n counts since the last event is simply the normalized survival function $w(n)/(\mu + 1)$. Appendix A derives the statistical complexity using this and Theorem 1. The resulting formulae are given in Table I for the various cases.

As described in Sec. II, we can also endow the causal state space with a transition dynamic in order to construct the renewal process ϵ-machine—the process's minimal unifilar hidden Markov model. The transition dy-
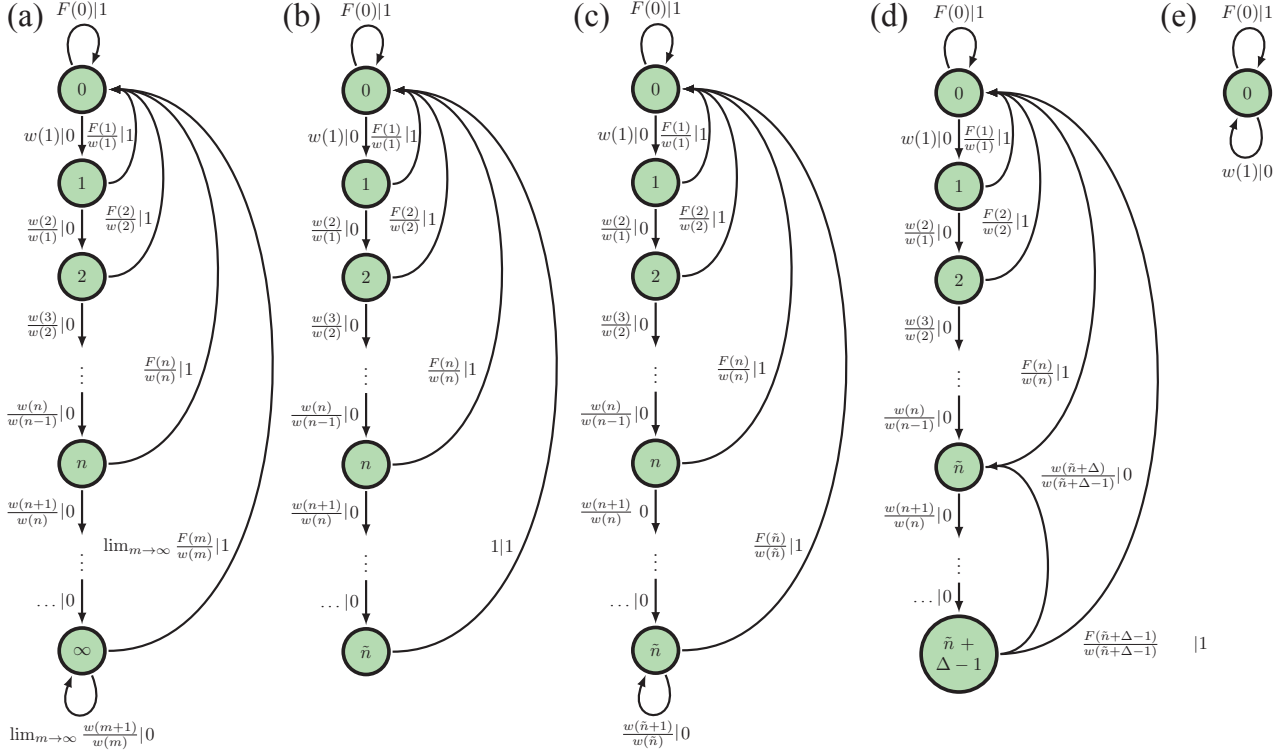
FIG. 3. ϵ-Machine architectures for discrete-time stationary renewal processes: (a) Not eventually Δ-Poisson with unbounded-support interevent distribution. (b) Not eventually Δ-Poisson with bounded-support interevent distribution. (c) Eventually Δ-Poisson with characteristic $(\tilde{n}, \Delta = 1)$ in Def. 2. (d) Eventually Δ-Poisson with characteristic $(\tilde{n}, \Delta > 1)$ in Def. 2. (e) For comparison, a Poisson process ϵ-machine.

namic is sensitive to $F(n)$'s support and not only its boundedness. For instance, the probability of observing an event given that it has been n counts since the last event is $F(n)/w(n)$. For the generic discrete-time renewal process this is exactly the transition probability from causal state n to causal state 0. If $F(n) = 0$, then there is no probability of transition from $\sigma = n$ to $\sigma = 0$. See App. A for details.

Figures 3(a)-3(d) display the causal state architectures, depicted as state-transition diagrams, for the ϵ-machines in the various cases delineated. Figure 3(a) is the ϵ-machine of a generic renewal process whose interevent interval can be arbitrarily large and whose interevent distribution never has exponential tails. Figure 3(b) is the ϵ-machine of a renewal process whose interevent distribution never has exponential tails but cannot have arbitrarily large interevent counts. The ϵ-machine in Fig. 3(c) looks quite similar to the ϵ-machine in Fig. 3(b), but it has an additional transition that connects the last state $\tilde{n}$ to itself. This added transition changes our structural interpretation of the process. Interevent counts can be arbitrarily large for this ϵ-machine but past an interevent count of $\tilde{n}$, the interevent distribution is exponential. Finally,

the ϵ-machine in Fig. 3(d) represents an eventually Δ-Poisson process with $\Delta > 1$ whose structure is conceptually most similar to that of the ϵ-machine in Fig. 3(c). (See Def. 2 for the precise version of that statement.) If our renewal process disallows seeing interevent counts of a particular length L, then this will be apparent from the ϵ-machine since there will be no transition between the causal state corresponding to an interevent count of L and causal state 0.

As described in Sec. II, we can analytically characterize a process' information architecture far better once we characterize its statistical structure in reverse time.

Lemma 2. *Groupings of futures with the same counts to next event are reverse-time prescient statistics for discrete-time stationary renewal processes.*

Theorem 2. *(a) The reverse-time causal states of a discrete-time stationary renewal process that is not eventually Δ-Poisson are groupings of futures with the same count to next event. (b) The reverse-time causal states of a discrete-time eventually Δ-Poisson stationary renewal process are groupings of futures with the same count to next event up until $\tilde{n}$ plus groupings of futures whose count since last event n are in the same equivalence class*

Quantity	Expression	
$C_\mu^+ = H[\mathcal{S}^+]$	$-\sum_{n=0}^\infty \frac{w(n)}{\mu+1} \log_2 \frac{w(n)}{\mu+1}$	Not eventually Δ-Poisson
	$-\sum_{n=0}^{\tilde{n}-1} \frac{w(n)}{\mu+1} \log_2 \frac{w(n)}{\mu+1} - \sum_{m=0}^{\Delta-1} \frac{\sum_{k=0}^\infty w(\tilde{n}+k\Delta+m)}{\mu+1} \log_2 \frac{\sum_{k=0}^\infty w(\tilde{n}+k\Delta+m)}{\mu+1}$	Eventually Δ-Poisson
$\mathbf{E} = I[X_{:0}; X_{0:}]$	$-2\sum_{n=0}^\infty \frac{w(n)}{\mu+1} \log_2 \frac{w(n)}{\mu+1} + \sum_{n=0}^\infty (n+1) \frac{F(n)}{\mu+1} \log_2 \frac{F(n)}{\mu+1}$	
$h_\mu = H[X_0|X_{:0}]$	$-\frac{1}{\mu+1} \sum_{n=0}^\infty F(n) \log_2 F(n)$	
$b_\mu = I[X_{1:}; X_0|X_{:0}]$	$\frac{1}{\mu+1} \left\{ \sum_{n=0}^\infty (n+1)F(n) \log_2 F(n) - \sum_{m,n=0}^\infty g(m,n) \log_2 g(m,n) \right\}$	
$\sigma_\mu = I[X_{1:}; X_{:0}|X_0]$	$\frac{1}{\mu+1} \left\{ \mu \log_2 \mu + \sum_{n=0}^\infty nF(n) \log_2 F(n) - 2\sum_{n=0}^\infty w(n) \log_2 w(n) \right\}$	
$q_\mu = I[X_{1:}; X_0; X_{:0}]$	$\frac{1}{\mu+1} \left\{ \sum_{m,n=0}^\infty g(m,n) \log_2 g(m,n) - \sum_{n=0}^\infty w(n) \log_2 w(n) + (\mu+1)\log_2(\mu+1) - \mu \log_2 \mu \right\}$	
$r_\mu = H[X_0|X_{1:}, X_{:0}]$	$\frac{1}{\mu+1} \left\{ \sum_{m,n=0}^\infty g(m,n) \log_2 g(m,n) - \sum_{n=0}^\infty (n+2)F(n) \log_2 F(n) \right\}$	
$H_0 = H[X_0]$	$-\frac{1}{\mu+1} \log_2 \frac{1}{\mu+1} - \left(1 - \frac{1}{\mu+1}\right) \log_2 \left(1 - \frac{1}{\mu+1}\right)$	

TABLE I. Structural measures and information architecture of a stationary renewal process with interevent counts drawn from the distribution $F(n)$, $n \geq 0$, survival count distribution $w(n) = \sum_{m=n}^\infty F(m)$, and mean interevent count $\mu = \sum_{n=0}^\infty nF(n) < \infty$. The function $g(m,n)$ is defined by $g(m,n) = F(m+n+1) + F(m)F(n)$. Cases are needed for C_μ but not other quantities, such as block entropy and information architecture quantities, since the latter can be calculated just as well from prescient machines. The quantities χ (crypticity) and $\mathbf{E}(M, N) = I[X_{-M:0}; X_{0:N}]$ are no less interesting than the others given here, but their expressions are not compact; see App. B.

as $\tilde{n}$ modulo Δ.

As a result, in reverse time a stationary renewal process is effectively the same stationary renewal process—counts between events are still independently drawn from $F(n)$. Thus, the causal irreversibility vanishes: $\Xi = 0$.

Moreover, these results taken together indicate that we can straightforwardly build a renewal process's *bidirectional machine* from these forward and reverse-time causal states, as described in Refs. [37, 38, 46]. Additional properties can then be deduced from the bidirectional machine, but we leave this for the future.

In closing, we note a parallel in Ref. [47]: Their A_t is the forward-time prescient statistic and E_t is the reverse-time prescient statistic. They call these the backward and forward recurrence times, respectively. They address continuous-time processes, however, and so the setting differs somewhat.

IV. INFORMATION ARCHITECTURE OF RENEWAL PROCESSES

As Sec. II described, many quantities that capture a process's predictability and randomness can be calculated from knowing the block entropy function $H(L)$. Often, the block entropy is estimated by generating samples of a process and estimating the entropy of a trajectory distribution. This method has the obvious disadvantage that at large L, there are $|\mathcal{A}|^L$ possible trajectories

and $|\mathcal{A}|^{h_\mu L}$ typical trajectories. And so, one easily runs into the problem of severe undersampling, previously referred to as the curse of dimensionality. This matters most when the underlying process has long-range temporal correlations.

Nor can one calculate the block entropy and other such information measures exactly from generative models that are not maximally predictive (prescient). Then, the model states do not shield the past from the future. For instance, as noted above, one cannot calculate the SNS's entropy rate from its simple two-state generative HMM. The entropy of the next symbol given the generative model's current state (A or B) actually *underestimates* the true entropy rate by assuming that we can almost always precisely determine the underlying model state from the past. For a sense of the fundamental challenge, see Refs. [25, 48].

However, we *can* calculate the block entropy and various other information measures in closed-form from a maximally predictive model. In other words, finding an ϵ-machine allows one to avoid the curse of dimensionality inherently involved in calculating the entropy rate, excess entropy, or the other information measures discussed here.

Figure 1 summarized the above points. This section makes good on the procedure outlined there by providing analytic formulae for various information measures of renewal processes. The formulae for the entropy rate of

a renewal process is already well known, but all others are new.

Prescient HMMs built from the prescient statistics of Lemma 1 are maximally predictive models, and correspond to the unifilar Hidden Markov Model shown in Fig. 3(a). The prescient machines make no distinction between eventually Δ-Poisson renewal processes and one that is not, but they do contain information about the support of $F(n)$ through their transition dynamics. (See App. A.) Appendix B describes how a prescient machine can be used to calculate all information architecture quantities—r_μ, b_μ, σ_μ, q_μ, and the more familiar Shannon entropy rate h_μ and excess entropy $\mathbf{E}$. A general strategy for calculating these quantities, as described in Sec. II and Refs. [19, 40], is to calculate b_μ, h_μ, $\mathbf{E}$, and $H[X_0]$, and then to derive the other quantities using the information-theoretic identities given in Sec. II.

Table I gives the results of these calculations. It helps one's interpretation to consider two base cases. For a Poisson process, we gain no predictive power by remembering specific pasts, and we would expect the statistical complexity, excess entropy, and bound information rate to vanish. The entropy rate and ephemeral information, though, are nonzero. One can check that this is, indeed, the case. For a periodic process with period T, in contrast, one can check that $\mu + 1 = T$, since the period is the length of the string of 0's (mean interevent interval μ) concatenated with the subsequent event $x = 1$. The statistical complexity and excess entropy of this process are $\log_2 T$ and the entropy rate is $h_\mu = 0$, as expected.

Calculating the predictable information $\mathbf{E}(M, N)$ requires identifying finite-time prescient statistics, since the predictable information is the mutual information between forward-time causal states over pasts of length M and reverse-time causal states over futures of length N. Such finite-time prescient statistics are identified in Corollary 1, below, and the predictable information is derived in App. B. The final expression is not included in Table I due to its length.

Corollary 1. *Forward-time (and reverse-time) finite-time M prescient states of a discrete-time stationary renewal process are the counts from (and to) the next event up until and including M.*

All of these quantities can be calculated using a mixed-state presentation, as described in Ref. [46], though the formulae developed there are as yet unable to describe processes with a countably infinite set of mixed states. Calculations of finite-time entropy rate estimates using a mixed-state presentation are consistent with all other results here, though. Purely for simplicity, we avoid discussing mixed-state presentations, though interested readers can find relevant definitions at Ref. [46].
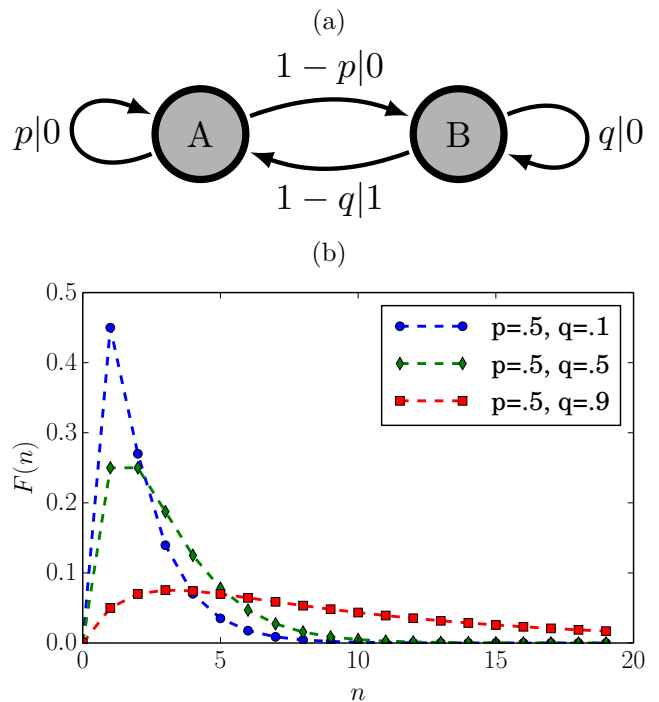

FIG. 4. (a) A (nonunifilar) hidden Markov model for the (p, q) parametrized SNS. (b) Example interevent distributions $F(n)$ from Eq. (15) for three parameter settings of (p, q).

V. NONUNIFILAR HMMS AND RENEWAL PROCESSES

The task of inferring an ϵ-machine for discrete-time, discrete-alphabet processes is essentially that of inferring minimal unifilar HMMs; what are sometimes also called "probabilistic deterministic" finite automata. In unifilar HMMs, the transition to the next hidden state given the previous one and next emitted symbol is determined. Nonunifilar HMMs are a more general class of time series models in which the transitions between underlying states given the next emitted symbol can be stochastic.

This simple difference in HMM structure has important consequences for calculating the predictable information, information architecture, and statistical complexity of time series generated by nonunifilar HMMs. First, note that for processes with a finite number of transient and recurrent causal states, these quantities can be calculated in closed form [24]. Second, the autocorrelation function and power spectrum can also be calculated in closed form for nonunifilar presentations [49]. Unlike these cases, though, most of Table I's quantities defy current calculational techniques. As a result, exact calculations of these prediction-related information measures for even the simplest nonunifilar HMMs can be surprisingly difficult.

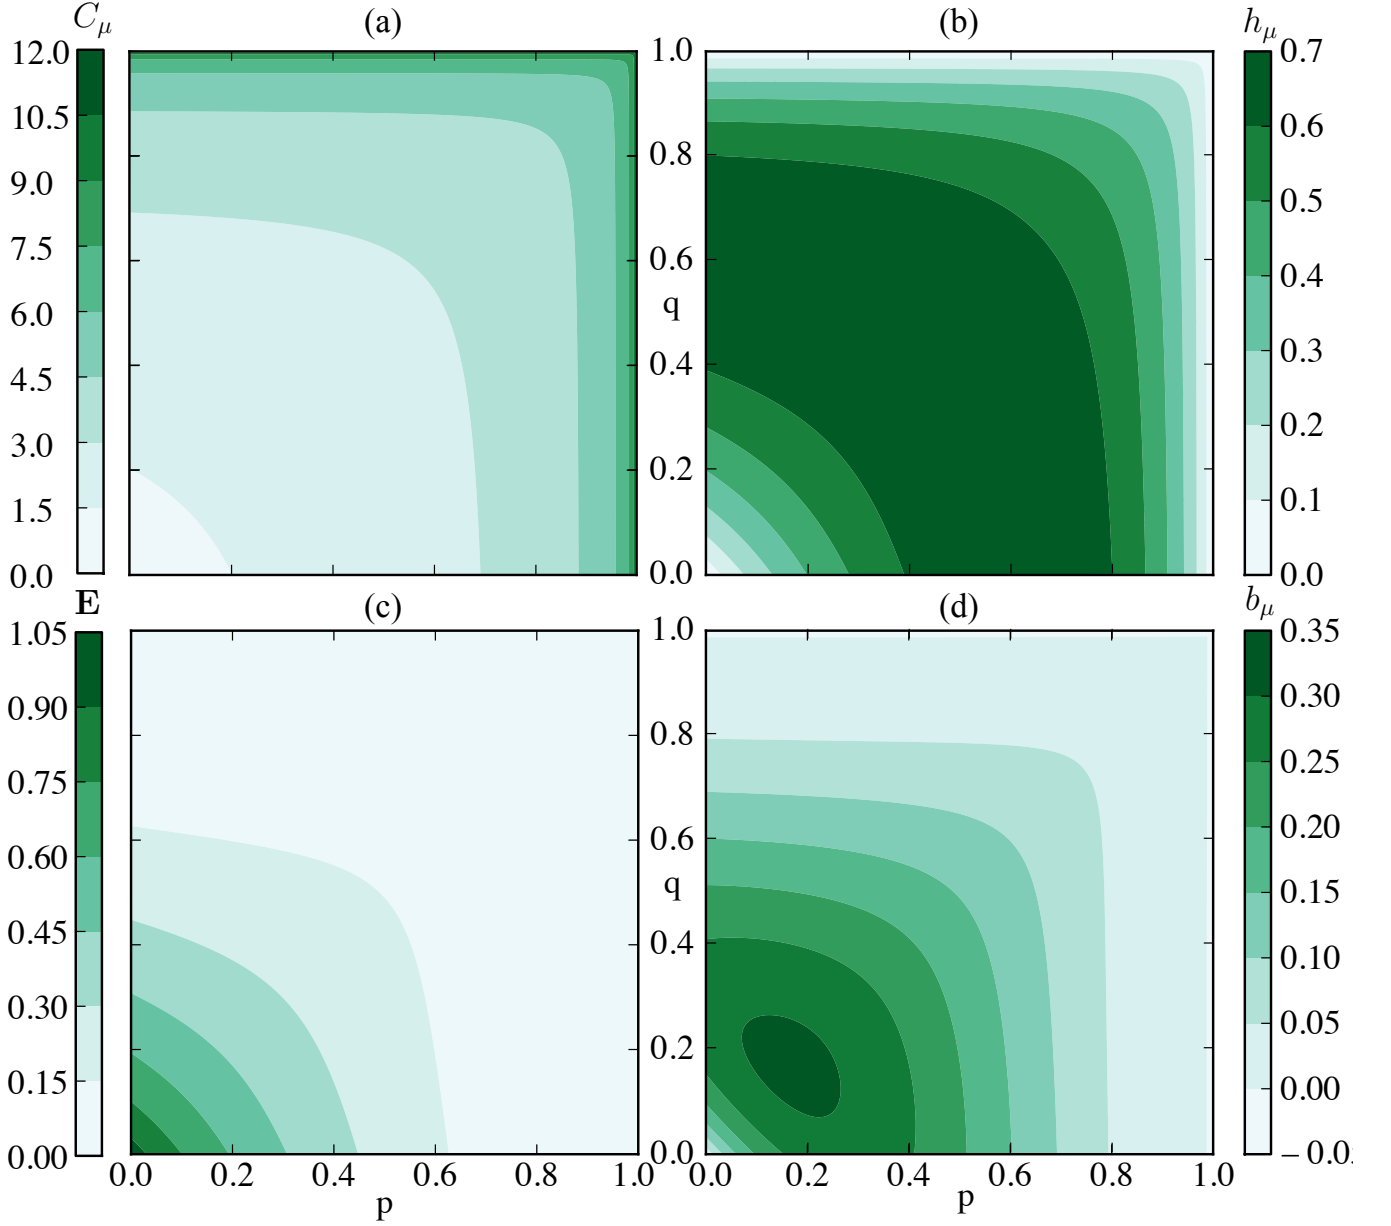To illustrate this point, we focus our attention on a

FIG. 5. Contour plots of various information measures (in bits) as functions of SNS parameters p and q. (a) C_μ, increasing when $F(n)$ has slower decay. (b) h_μ, higher when transition probabilities are maximally stochastic. (c) $\mathbf{E}$, higher the closer the SNS comes to period-2. (d) b_μ, highest between the maximally stochastic transition probabilities that maximize h_μ and maximally deterministic transition probabilities that maximize $\mathbf{E}$.

parametrized version of the SNS shown in Fig. 4(a). As for the original SNS in Fig. 2, transitions from state B are unifilar, but transitions from state A are not. As noted before, the time series generated by the parametrized SNS is a discrete-time renewal process with interevent

count distribution:

$$F(n) = \begin{cases} (1-p)(1-q)(p^n - q^n)/(p-q) & p \neq q \ , \\ (1-p)^2 n p^{n-1} & p = q \ . \end{cases}$$
(15)

Figure 4(b) also shows $F(n)$ at various parameter choices. The nonunifilar HMM there should be contrasted with the unifilar HMM presentation of the parametrized SNS

which is the ϵ-machine in Fig. 3(a), with a countable infinity of causal states.

Both parametrized SNS presentations are "minimally complex", but according to different metrics. On the one hand, the nonunifilar presentation is a minimal *generative* model: No one-state HMM (i.e., biased coin) can produce a time series with the same statistics. On the other, the unifilar HMM is the minimal *maximally predictive* model: In order to predict the future as well as possible given the entire past, we must at least remember how many 0's have been seen since the last 1. That memory requires a countable infinity of prescient states. The preferred complexity metric is a matter of taste and desired implementation, modulo important concerns regarding overfitting or ease of inference [26]. However, if we wish to calculate the information measures in Table I as accurately as possible, finding a maximally predictive model—a unifilar presentation, that is—is necessary.

The SNS is not eventually Δ-Poisson with an unbounded-support interevent distribution. It can only be approximated by an eventually Δ-Poisson process. And so, one estimates its information anatomy via infinite sums. Using the formulae of Table I, Fig. 5 shows how the statistical complexity C_μ, excess entropy $\mathbf{E}$, entropy rate h_μ, and bound information b_μ vary with the transition probabilities p and q. C_μ often reveals detailed information about a process' underlying structure, but for the parametrized SNS and other renewal processes, the statistical complexity merely reflects the spread of the interevent distribution. Thus, it increases with increasing p and q. $\mathbf{E}$, a measure of how much can be predicted rather than historical memory required for prediction, increases as p and q decrease. The intuition for this is that as p and q vanish, the process arrives at a perfectly predictable period-2 sequence. We see that the SNS constitutes a simple example of a class of processes over which information transmission between the past and future ($\mathbf{E}$) and information storage (C_μ) are anticorrelated. The entropy rate h_μ in Fig. 5(b) is maximized when transitions are uniformly stochastic and the bound information b_μ in Fig. 5(d) is maximized somewhere between fully stochastic and fully deterministic regimes.

Figure 6 presents a more nuanced decomposition of the information measures as $p = q$ vary from 0 to 1. Figure 6(a) breaks down the single-measurement entropy $H[X_0]$ into redundant information ρ_μ in a single measurement, predictively useless generated information r_μ, and predictively useful generated entropy b_μ. As p increases, the SNS moves from mostly predictable (close to period-2) to mostly unpredictable, shown by the relative height of the (green) line denoting h_μ to the (red) line denoting $H[X_0]$. The portion b_μ of h_μ predictive of the future is maximized at lower p when the single-measurement en-
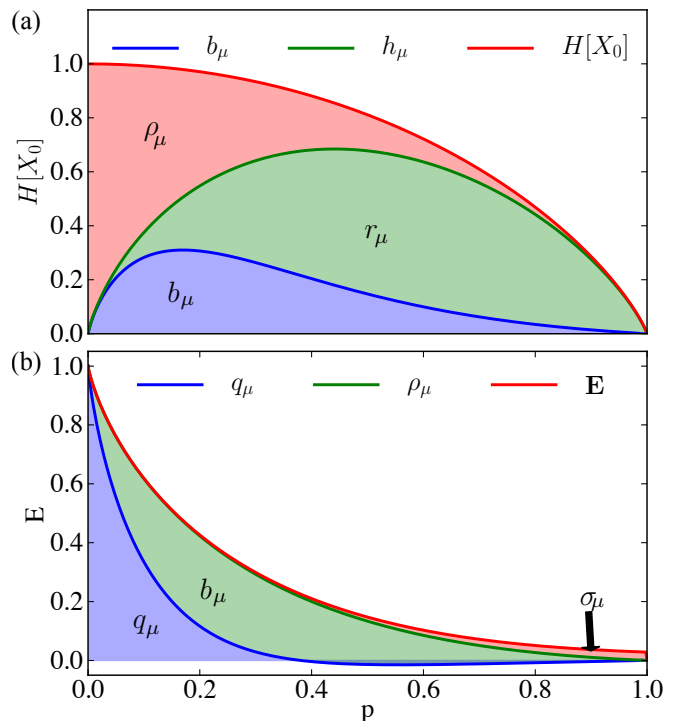

FIG. 6. (a) Simple Nonunifilar Source information architecture as a function of p with parameters $q = p$. The single-measurement entropy $H[X_0]$ is the upper solid (red) line, entropy rate h_μ the middle solid (green) line, the bound information b_μ the lower solid (blue) line. Thus, the blue area corresponds to b_μ, the green area to the ephemeral information $r_\mu = h_\mu - b_\mu$, and the red area to the single-symbol redundancy $\rho_\mu = H[X_0] - h_\mu$. (b) The components of the predictable information—the excess entropy $\mathbf{E} = \sigma_\mu + b_\mu + q_\mu$ in bits—also as a function of p with $q = p$. The lowest (blue) line is q_μ; the middle (green) line is $q_\mu + b_\mu$, so that the green area denotes b_μ's contribution to $\mathbf{E}$. The upper (red) line is $\mathbf{E}$, so that the red area denotes elusive information σ_μ in $\mathbf{E}$. Note that for a large range of p the co-information q_μ is (slightly) negative.

tropy is close to a less noisy period-2 process. Figure 6(b) decomposes the predictable information $\mathbf{E}$ into the predictable information hidden from the present σ_μ, the predictable generated entropy in the present b_μ, and the co-information q_μ shared between past, present, and future. Recall that the co-information $q_\mu = \mathbf{E} - \sigma_\mu - b_\mu$ can be negative and, for a large range of values, it is. Most of the predictable information passes through the present as indicated by σ_μ being a small for most parameters p. Hence, even though the parametrized SNS is technically an infinite-order Markov process, it can be well approximated by a finite-order Markov process without much predictable information loss, as noted previously with rate-distortion theory [50].

VI. CONCLUSIONS

Stationary renewal processes are well studied, easy to define, and, in many ways, temporally simple. Given this simplicity and their long history it is somewhat surprisingly that one is still able to discover new properties; in our case, by viewing them through an information-theoretic lens. Indeed, their simplicity becomes apparent in the informational and structural analyses. For instance, renewal processes are causally reversible with isomorphic ϵ-machines in forward and reverse-time, i.e., temporally reversible. Applying the causal-state equivalence relation to renewal processes, however, also revealed several unanticipated subtleties. For instance, we had to delineate new subclasses of renewal process ("eventually Δ-Poisson") in order to completely classify ϵ-machines of renewal processes. Additionally, the informational architecture formulae in Table I are surprisingly complicated, since exactly calculating these informational measures requires a unifilar presentation. In Sec. V, we needed an infinite-state machine to study the informational architecture of a process generated by simple two-state HMM.

Looking to the future, the new structural view of renewal processes will help improve inference methods for infinite-state processes, as it tells us what to expect in the ideal setting—what are the effective states, what are appropriate null models, how informational quantities scale, and the like. For example, Figs. 3(a)-3(e) gave all possible causal architectures for discrete-time stationary renewal processes. Such a classification will allow for more efficient Bayesian inference of ϵ-machines of point processes, as developed in Ref. [26]. That is, we can leverage "expert" knowledge that one is seeing a renewal process to delineate the appropriate subset of model architectures and thereby avoid searching over the superexponentially large set of all HMM topologies.

The range of the results' application is much larger than that explicitly considered here. The formulae in Table I will be most useful for understanding renewal processes with infinite statistical complexity. For instance, Ref. [28] applies the formulae to study the divergence of the statistical complexity of continuous-time processes as the observation time scale decreases. And, Ref. [29] applies these formulae to renewal processes with infinite excess entropy. In particular, there we investigate the causal architectures of infinite-state processes that generate so-called *critical phenomena*—behavior with power-law temporal or spatial correlations [51]. The analysis of such critical systems often turns on having an appropriate *order parameter*. The statistical complexity and excess entropy are application-agnostic order parameters [52–54] that allow one to better quantify when a phase transition in stochastic processes has or has not occurred, as seen in Ref. [29]. Such critical behavior has even been implicated in early studies of human communication [55] (though see Refs. [56, 57]) and recently in neural dynamics [58] and in socially constructed, communal knowledge systems [59].

Appendix A: Causal Architecture

Notation. *Rather than write pasts and futures as semi-infinite sequences, we notate a past as a list of nonnegative integers [60]. The semi-infinite past $X_{:0}$ is equivalent to a list of interevent counts $N_{:0}$ and the count N_0' since last event. Similarly, the semi-infinite future $X_{0:}$ is equivalent to the count to next event $N_0 - N_0'$ and future interevent counts $N_{1:}$.*

Now, recall Lemma 1.

Lemma 1 *The counts since last event are prescient statistics of a discrete-time stationary renewal process.*

Proof. *This follows almost immediately from the definition of stationary renewal process and the definition of causal states, since the random variables N_i are all i.i.d.. Then:*

$$\Pr(X_{0:}|X_{:0}) = \Pr(N_0 - N_0'|N_0') \prod_{i=1}^{\infty} \Pr(N_i) \ .$$

And, therefore, $\Pr(X_{0:}|X_{:0} = x_{:0}) = \Pr(X_{0:}|X_{:0} = x_{:0}')$ is equivalent to $\Pr(N_0 - N_0'|N_0' = n_0) = \Pr(N_0 - N_0'|N_0' = n_0')$. Hence, the counts since last event are prescient.

In light of Lemma 1, we introduce new notation to efficiently refer to groups of pasts with the same count since last event.

Notation. *Let $r_n^+ := \{\overleftarrow{x} : x_{-n-1:0} = 10^n\}$ for $n \in \mathbb{Z}_{\geq 0}$. Recall that $10^n = 100 \cdots 00$, the sequence with n $0s$ following a 1.*

Remark. *Note that $\mathcal{R}^+ = \{r_n^+\}_{n=0}^{\infty}$ is always at least a forward-time prescient rival, if not the forward-time*

causal states $\boldsymbol{\mathcal{S}}^+$. The probability distribution over r_n^+ is straightforward to derive. Saying that $N_0' = n$ means there were n 0s since the last event, so that the symbol at X_{-n-1} must have been a 1. That is:

$$
\begin{aligned}
\pi(r_n^+) &= \Pr(N_0' = n) \\
&= \sum_{x \in \mathcal{A}} \Pr(N_0' = n, X_{-n-1} = x) \\
&= \Pr(N_0' = n, X_{-n-1} = 1) \\
&= \Pr(X_{-n-1} = 1) \Pr(X_{-n:0} = 0^n | X_{-n-1} = 1) \ .
\end{aligned}
$$

Since this is a stationary process, $\Pr(X_{-n-1} = 1)$ is independent of n, implying:

$$
\begin{aligned}
\pi(r_n^+) &\propto \Pr(X_{-n:0} = 0^n | X_{-n-1} = 1) \\
&= \sum_{m=0}^{\infty} \Pr(X_{-n:m+1} = 0^{n+m} 1 | X_{-n-1} = 1) \\
&= \sum_{m=n}^{\infty} F(m) \\
&= w(n) \ .
\end{aligned}
$$

We see that $\pi(r_n^+) = w(n)/Z$, with Z a normalization constant that makes $\sum_{n=0}^{\infty} w(n) = \mu + 1$. And so:

$$
\pi(r_n^+) = \frac{w(n)}{\mu + 1} \ .
$$

In the main text, Thm. 1 was stated with less precision so as to be comprehensible. Here, we state it with more precision, even though the meaning is obfuscated somewhat by doing so. In the proof, we still err somewhat on the side of comprehensibility, and so one might view this proof as more of a proof sketch.

Theorem 1 *The forward-time causal states of a discrete-time stationary renewal process that is not eventually Δ-Poisson are exactly $\boldsymbol{\mathcal{S}}^+ = \boldsymbol{\mathcal{R}}^+$, if F has unbounded support. When the support is bounded such that $F(n) = 0$ for all $n \geq N$, $\boldsymbol{\mathcal{S}}^+ = \{r_n^+\}_{n=0}^{N}$. Finally, a discrete-time eventually Δ-Poisson renewal process with characteristic $(\tilde{n}, \Delta)$ has forward-time causal states:*

$$
\boldsymbol{\mathcal{S}}^+ = \{r_n^+\}_{n=0}^{\tilde{n}-1} \cup \{\cup_{k=0}^{\infty} r_{\tilde{n}+k\Delta+m}^+\}_{m=0}^{\Delta-1} \ .
$$

This is a complete classification of the causal states of any persistent renewal process.

Proof. *From the proof of Lemma 1 in this appendix, we know that two prescient states r_n^+ and $r_{n'}^+$ are minimal only when:*

$$
\Pr(N_0 - N_0' | N_0' = n) = \Pr(N_0 - N_0' | N_0' = n') \ . \quad \text{(A1)}
$$

Since $\Pr(N_0 - N_0' = m | N_0' = n) = \Pr(N_0 = m + n)/\Pr(N_0' = n)$, $\Pr(N_0 = m + n) = F(m + n)$, and $\Pr(N_0' = n) = w(n)/(\mu+1)$ from earlier, we find that the equivalence class condition becomes:

$$
\frac{F(m + n)}{w(n)} = \frac{F(m + n')}{w(n')} \ , \quad \text{(A2)}
$$

for all $m \geq 0$.

First, note that for these conditional probabilities even to be well defined, $w(n) > 0$ and $w(n') > 0$. Hence, if F has bounded support—$\max \mathrm{supp} F(n) = N$—then the causal states do not include any r_n^+ for $n > N$. Furthermore, Eq. (A2) cannot be true for all $m \geq 0$, unless $n = n'$ for n and $n' \leq N$. To see this, suppose that $n \neq n'$ but that Eq. (A2) holds. Then choose $m = N+1-\max(n, n')$ to give $0 = F(N + 1 - |n - n'|)/w(n')$, a contradiction unless $n = n'$.

So, for all remaining cases, we can assume that F in Eq. (A2) has unbounded support.

A little rewriting makes the connection between Eq. (A2) and an eventually Δ-Poisson process clearer. First, we choose $m = 0$ to find:

$$
\frac{F(n)}{w(n)} = \frac{F(n')}{w(n')},
$$

which we can use to rewrite Eq. (A2) as:

$$
\frac{F(m + n)}{F(n)} = \frac{F(m + n')}{F(n')},
$$

or more usefully:

$$
F(n' + m) = \frac{F(n')}{F(n)} F(n + m).
$$

A particularly compact way of rewriting this is to define $\Delta' := n' - n$, which gives $F(n' + m) = F((n + m) + \Delta')$. In this form, it is clear that the above equation is a recurrence relation on F in steps of Δ', so that we can write:

$$
F((n + m) + k\Delta') = \left(\frac{F(n')}{F(n)}\right)^k F(n + m) \ . \quad \text{(A3)}
$$

This must be true for every $m \geq 0$. Importantly, since $w(n) = \sum_{m=n}^{\infty} F(m)$, satisfying this recurrence relation is equivalent to satisfying Eq. (A2). But Eq. (A3) is just the definition of an eventually Δ-Poisson process in disguise; relabel with $\lambda := F(n')/F(n)$, $\tilde{n} := n$, and $\Delta = \Delta'$.

Therefore, if Eq. (A2) does not hold for any pair $n \neq n'$, the process is not eventually Δ-Poisson and the prescient states identified in Lemma 1 are minimal; i.e., they are the causal states.

If Eq. (A2) does hold for some $n \neq n'$, choose the minimal such n and n' both. The renewal process is eventually Δ-Poisson with characterization $\Delta = n' - n$ and $\tilde{n}$. And, $F(\tilde{n} + m)/w(\tilde{n} + m) = F(\tilde{n} + m')/w(\tilde{n} + m')$ implies that $m \equiv m' \mod \Delta$ since otherwise, the n and n' chosen would not be minimal. Hence, the causal states are exactly those given in the theorem's statement.

Remark. *For the resulting $F(n)$ to be a valid interevent distribution, $\lambda = F(\tilde{n} + \Delta)/F(\tilde{n}) < 1$ as normalization implies:*

$$\sum_{n=0}^{\tilde{n}-1} F(n) + \sum_{n=\tilde{n}}^{\tilde{n}+\Delta-1} \frac{F(n)}{1-\lambda} = 1 \ .$$

Notation. *Let's denote $\boldsymbol{\mathcal{S}}^+ = \{\sigma_n^+ := r_n^+\}_{n=0}^\infty$ for a renewal process that is not eventually Δ-Poisson, $\boldsymbol{\mathcal{S}}^+ = \{\sigma_n^+ := r_n^+\}_{n=0}^{\tilde{n}}$ for an eventually Δ-Poisson renewal process with bounded support, and $\boldsymbol{\mathcal{S}}^+ = \{\sigma_n^+ := r_n^+\}_{n=0}^{\tilde{n}-1} \cup \{\sigma_{\tilde{n}+m}^+ := \cup_{k=0}^\infty r_{\tilde{n}+k\Delta+m}^+\}_{m=0}^{\Delta-1}$ for an eventually Δ-Poisson process.*

The probability distribution over these forward-time causal states is straightforward to derive from $\pi(r_n^+) = w(n)/(\mu+1)$. So, for a renewal process that is not eventually Δ-Poisson or one that is with bounded support, $\pi(\sigma_n^+) = w(n)/(\mu+1)$. (For the latter, n only runs from 0 to $\tilde{n}$.) For an eventually Δ-Poisson renewal process $\pi(\sigma_n^+) = w(n)/(\mu+1)$ when $n < \tilde{n}$ and:

$$\pi(\sigma_n^+) = \sum_{k=0}^\infty \pi(r_{n+k\Delta}^+)$$
$$= \frac{\sum_{k=0}^\infty w(n+k\Delta)}{\mu+1} \ ,$$

when $\tilde{n} \leq n < \tilde{n} + \Delta$. And so, the statistical complexity given in Table I follows from $C_\mu^+ = H[\mathcal{S}^+]$.

Recall Lemma 2 and Thm. 2.

Lemma 2 *Groupings of futures with the same counts to next event are reverse-time prescient statistics for discrete-time stationary renewal processes.*

Theorem 2 *(a) The reverse-time causal states of a discrete-time stationary renewal process that is not eventually Δ-Poisson are groupings of futures with the same count to next event up until and including N, if N is finite. (b) The reverse-time causal states of a discrete-time eventually Δ-Poisson stationary renewal process are groupings of futures with the same count to next event up until $\tilde{n}$, plus groupings of futures whose count since last event n are in the same equivalence class as $\tilde{n}$ modulo Δ.*

Proof. *The proof for both claims relies on a single fact: In reverse-time, a stationary renewal process is still a stationary renewal process with the same interevent count distribution. The lemma and theorem therefore follow from Lemma 1 and Thm. 1.*

Since the forward and reverse-time causal states are the same with the same future conditional probability distribution, we have $C_\mu^+ = C_\mu^-$ and the causal irreversibility vanishes: $\Xi = 0$.

Transition probabilities can be derived for both the renewal process's prescient states and its ϵ-machine as follows. For the prescient machine, if a 0 is observed when in r_n^+, we transition to r_{n+1}^+; else, we transition to r_0^+ since we just saw an event. Basic calculations show that these transition probabilities are:

$$T_{r_n^+ r_m^+}^{(x)} = \Pr(\mathcal{R}_{t+1}^+ = r_m^+, X_{t+1} = x | \mathcal{R}_t^+ = r_n^+)$$
$$= \frac{F(n)}{w(n)} \delta_{m,0} \delta_{x,1} + \frac{w(n+1)}{w(n)} \delta_{m,n+1} \times \delta_{x,0} \ .$$

Not only do these specify the prescient machine transition dynamic but, due to the close correspondence between prescient and causal states, they also automatically give the ϵ-machine transition dynamic:

$$T_{\sigma\sigma'}^{(x)} = \Pr(\mathcal{S}_{t+1}^+ = \sigma', X_{t+1} = x | \mathcal{S}_t^+ = \sigma)$$
$$= \sum_{r,r' \in \boldsymbol{\mathcal{R}}^+} T_{r' \to r}^{(x)} \Pr(\mathcal{S}_{t+1}^+ = \sigma' | \mathcal{R}_{t+1}^+ = r)$$
$$\times \Pr(\mathcal{R}_t^+ = r' | \mathcal{S}_t^+ = \sigma) \ .$$

Appendix B: Information Architecture

It is straightforward to show that $\Pr(X_0 = 0) = \frac{1}{\mu+1}$ and, thus:

$$H[X_0] = -\frac{1}{\mu+1} \log_2 \frac{1}{\mu+1}$$
$$- \left(1 - \frac{1}{\mu+1}\right) \log_2 \left(1 - \frac{1}{\mu+1}\right) \ .$$

The entropy rate is readily calculated from the prescient machine:

$$h_\mu = \sum_{n=0}^\infty H[X_{t+1} | \mathcal{R}_t^+ = r_n^+] \pi(r_n^+)$$
$$= -\sum_{n=0}^\infty \frac{w(n)}{\mu+1} \left(\frac{F(n)}{w(n)} \log_2 \frac{F(n)}{w(n)}\right.$$
$$\left. + \frac{w(n+1)}{w(n)} \log_2 \frac{w(n+1)}{w(n)}\right) \ .$$

And, after some algebra, this simplifies to:

$$h_\mu = -\frac{1}{\mu+1} \sum_{n=0}^{\infty} F(n) \log_2 F(n) \ ,$$

once we recognize that $w(0) = 1$ and so $w(0) \log_2 w(0) = 0$ and we recall that $w(n+1) + F(n) = w(n)$. The excess entropy, being the mutual information between forward and reverse-time prescient states is [37, 46]:

$$\begin{aligned}
\mathbf{E} &= I[\mathcal{R}^+; \mathcal{R}^-] \\
&= H[\mathcal{R}^+] - H[\mathcal{R}^+|\mathcal{R}^-] \ .
\end{aligned}$$

And so, to calculate, we note that:

$$\Pr(r_n^+, r_m^-) = \frac{F(m+n)}{\mu+1} \text{ and}$$
$$\Pr(r_n^+|r_m^-) = \frac{F(n+m)}{w(m)} \ .$$

After some algebra, we find that:

$$H[\mathcal{R}^+] = -\sum_{n=0}^{\infty} \frac{w(n)}{\mu+1} \log_2 \frac{w(n)}{\mu+1}$$

and that:

$$\begin{aligned}
H[\mathcal{R}^+|\mathcal{R}^-] &= -\sum_{m,n=0}^{\infty} \frac{F(n+m)}{\mu+1} \log_2 \frac{F(n+m)}{w(m)} \\
&= -\sum_{m=0}^{\infty} \frac{m+1}{\mu+1} F(m) \log_2 \frac{F(m)}{\mu+1} \\
&\quad + \sum_{m=0}^{\infty} \frac{w(m)}{\mu+1} \log_2 \frac{w(m)}{\mu+1} \ .
\end{aligned}$$

The above quantity is the forward crypticity χ^+ [37] when the renewal process is not eventually Δ-Poisson. These together imply:

$$\begin{aligned}
\mathbf{E} &= -2 \sum_{n=0}^{\infty} \frac{w(n)}{\mu+1} \log_2 \frac{w(n)}{\mu+1} \\
&\quad + \sum_{m=0}^{\infty} (m+1) \frac{F(m)}{\mu+1} \log_2 \frac{F(m)}{\mu+1} \ .
\end{aligned}$$

And, finally, the bound information b_μ is:

$$\begin{aligned}
b_\mu &= I[X_{1:}; X_0 | X_{:0}] \\
&= I[\mathcal{R}_1^-; X_0 | \mathcal{R}_0^+] \\
&= H[\mathcal{R}_1^-|\mathcal{R}_0^+] - H[\mathcal{R}_1^-|\mathcal{R}_1^+] \ ,
\end{aligned}$$

where we used the causal shielding properties of prescient states, $X_{:0} \to \mathcal{R}_0^+ \to \mathcal{R}_1^- \to X_{1:}$, and the unifilarity

of the prescient machines as shown in Figs. 3(a)-3(d). While we already calculated $H[\mathcal{R}_1^-|\mathcal{R}_1^+]$, we still need to calculate $H[\mathcal{R}_1^-|\mathcal{R}_0^+]$. We do so using the prescient machine's transition dynamic. In particular:

$$\begin{aligned}
&\Pr(\mathcal{R}_1^- = n | \mathcal{R}_0^+ = m) \\
&= \sum_{r \in \mathcal{R}^+} \Pr(\mathcal{R}_1^- = n | \mathcal{R}_1^+ = r) \Pr(\mathcal{R}_1^+ = r | \mathcal{R}_0^+ = m) \\
&= \frac{F(m+n+1) + F(n)F(m)}{w(m)} \ .
\end{aligned}$$

Where we omit details getting to the last line. Eventually, the calculation yields:

$$\begin{aligned}
b_\mu &= \frac{\sum_{n=0}^{\infty}(n+1)F(n)\log_2 F(n)}{\mu+1} \\
&\quad - \frac{\sum_{m,n=0}^{\infty} g(m,n) \log_2 g(m,n)}{\mu+1} \ ,
\end{aligned}$$

where:

$$g(m,n) = F(m+n+1) + F(n)F(m) \ .$$

From the expressions above, we immediately solve for $r_\mu = h_\mu - b_\mu$, $q_\mu = H[X_0] - h_\mu - b_\mu$, and $\sigma_\mu = \mathbf{E} - q_\mu$. Thereby laying out information architecture of stationary renewal processes.

Finally, we calculate the finite-time predictable information $\mathbf{E}(M, N)$ as the mutual information between finite-time forward and reverse-time prescient states:

$$\mathbf{E}(M, N) = H[\mathcal{R}^{-N}] - H[\mathcal{R}^{-N}|\mathcal{R}^{+M}] \ . \tag{B1}$$

Recall Corollary 1.

Corollary 1 *Forward-time (and reverse-time) finite-time M prescient states of a discrete-time stationary renewal process are the counts from (and to) the next event up until and including M.*

Proof. *From Lemmas 1 and 2, we know that counts from (to) the last (next) event are prescient forward-time (reverse-time) statistics. If our window on pasts (futures) is M, then we cannot distinguish between counts since (to) the last (next) event that are M and larger. Hence, the finite-time M prescient statistics are the counts from (and to) the next event up until and including M, where a finite-time M prescient state includes all pasts with M or more counts from (to) the last (next) event.*

To calculate $\mathbf{E}(M, N)$, we find $\Pr(\mathcal{R}^{+M}, \mathcal{R}^{-N})$ by marginalizing $\Pr(\mathcal{R}^+, \mathcal{R}^-)$. For ease of notation, we first define a function:

$$u(m) = \sum_{n=m}^{\infty} w(n) \ .$$

Algebra not shown here yields:

$$\mathbf{E}(M, N) = H[\mathcal{S}^{-N}] - H[\mathcal{S}^{-N}|\mathcal{S}^{+M}]$$

$$= \log_2(\mu + 1) - \frac{\sum_{n=0}^{N-1} w(n) \log_2 w(n)}{\mu + 1}$$

$$- \frac{\sum_{m=0}^{M-1} w(m) \log_2 w(m)}{\mu + 1}$$

$$+ \frac{\sum_{n=M}^{N+M-1} w(n) \log_2 w(n)}{\mu + 1}$$

$$+ \frac{\sum_{n=N}^{N+M-1} w(n) \log_2 w(n)}{\mu + 1}$$

$$- \frac{u(N) \log_2 u(N) + u(M) \log_2 u(M)}{\mu + 1}$$

$$+ \frac{u(N + M) \log_2 u(N + M)}{\mu + 1}$$

$$+ \frac{\sum_{m=0}^{M-1} \sum_{n=m}^{N+m-1} F(n) \log_2 F(n)}{\mu + 1} .$$

Two cases of interest are equal windows ($N = M$) and semi-infinite pasts ($M \to \infty$). In the former, we find:

$$\mathbf{E}(M, M) = \log_2(\mu + 1) - \frac{2 \sum_{m=0}^{M-1} w(m) \log_2 w(m)}{\mu + 1}$$

$$+ \frac{2 \sum_{m=M}^{2M-1} w(m) \log_2 w(m)}{\mu + 1}$$

$$- \frac{2u(M) \log_2 u(M)}{\mu + 1} + \frac{u(2M) \log_2 u(2M)}{\mu + 1}$$

$$+ \frac{\sum_{m=0}^{M-1} \sum_{n=m}^{M+m-1} F(n) \log_2 F(n)}{\mu + 1} .$$

In the latter case of semi-infinite pasts several terms vanish and we have:

$$\mathbf{E}(N) = \log_2(\mu + 1) - \frac{2 \sum_{n=0}^{N-1} w(n) \log_2 w(n)}{\mu + 1}$$

$$- \frac{u(N) \log_2 u(N)}{\mu + 1}$$

$$+ \frac{N \sum_{n=N}^{\infty} F(n) \log_2 F(n)}{\mu + 1}$$

$$+ \frac{\sum_{n=0}^{N-1} (n + 1) F(n) \log_2 F(n)}{\mu + 1} .$$

[1] Smith, W.L. Renewal Theory and Its Ramifications. *J. Roy. Stat. Soc. B* **1958**, *20*, 243–302.

[2] Gerstner, W.; Kistler. Statistics of spike trains. In *Spiking Neuron Models*; Cambridge University Press: Cambridge, United Kingdom, 2002.

[3] Beichelt, F. *Stochastic Processes in Science, Engineering and Finance*; Chapman and Hall: New York, 2006.

[4] Barbu, V.S.; Limnios, N. Semi-Markov chains and hidden semi-Markov models toward applications. *Their Use in Reliability and DNA Analysis* **2008**, *191*.

[5] Lowen, S.B.; Teich, M.C. Fractal Renewal Processes Generate 1/f Noise. *Phys. Rev. E* **1993**, *47*, 992–1001.

[6] Lowen, S.B.; Teich, M.C. Fractal Renewal Processes. *IEEE Trans. Info. Th.* **1993**, *39*, 1669–1671.

[7] Cakir, R.; Grigolini, P.; Krokhin, A.A. Dynamical Origin of Memory and Renewal. *Phys. Rev. E* **2006**, *74*, 021108.

[8] Akimoto, T.; Hasumi, T.; Aizawa, Y. Characterization of intermittency in renewal processes: Application to earthquakes. *Phys. Rev. E* **2010**, *81*, 031133.

[9] Montero, M.; Villarroel, J. Monotonic Continuous-time Random Walks with Drift and Stochastic Reset Events. *Phys. Rev. E* **2013**, *87*, 012116.

[10] Bologna, M.; West, B.J.; Grigolini, P. Renewal and Memory Origin of Anomalous Diffusion: A discussion of their joint action. *Phys. Rev. E* **2013**, *88*, 062106.

[11] Bianco, S.; Ignaccolo, M.; Rider, M.S.; Ross, M.J.; Winsor, P.; Grigolini, P. Brain, Music, and Non-Poisson Renewal Processes. *Phys. Rev. E* **2007**, *75*, 061911.

[12] Li, C.B.; Yang, H.; Komatsuzaki, T. Multiscale Complex Network of Protein Conformational Fluctuations in Single-Molecule Time Series. *Proc. Natl. Acad. Sci. USA* **2008**, *105*, 536–541.

[13] Kelly, D.; Dillingham, M.; Hudson, A.; Wiesner, K. A New Method for Inferring Hidden Markov Models from Noisy Time Sequences. *PLoS One* **2012**, *7*, e29703.

[14] Onaga, T.; Shinomoto, S. Bursting Transition in a Linear Self-exciting Point Process. *Phys. Rev. E* **2014**, *89*, 042817.

[15] Valenza, G.; Citi, L.; Scilingo, E.P.; Barbieri, R. Inhomogeneous Point-Process Entropy: An Instantaneous

measure of complexity in discrete systems. *Phys. Rev. E* **2014**, *89*, 052803.

[16] Shalizi, C.R.; Crutchfield, J.P. Information Bottlenecks, Causal States, and Statistical Relevance Bases: How to Represent Relevant Information in Memoryless Transduction. *Adv. Complex Sys.* **2002**, *5*, 91–95.

[17] Still, S.; Crutchfield, J.P. Structure or Noise? **2007**. Santa Fe Institute Working Paper 2007-08-020; arxiv.org physics.gen-ph/0708.0654.

[18] Still, S.; Crutchfield, J.P.; Ellison, C.J. Optimal Causal Inference: Estimating Stored Information and Approximating Causal Architecture. *CHAOS* **2010**, *20*, 037111.

[19] Marzen, S.; Crutchfield, J.P. Information Anatomy of Stochastic Equilibria. *Entropy* **2014**, *16*, 4713–4748.

[20] Martius, G.; Der, R.; Ay, N. Information driven self-organization of complex robotics behaviors. *PLoS One* **2013**, *8*, e63400.

[21] Crutchfield, J.P.; Feldman, D.P. Regularities Unseen, Randomness Observed: Levels of Entropy Convergence. *CHAOS* **2003**, *13*, 25–54.

[22] Debowski, L. On Hidden Markov Processes with Infinite Excess Entropy. *J. Theo. Prob.* **2012**, *27*, 539–551.

[23] Travers, N.; Crutchfield, J.P. Infinite Excess Entropy Processes with Countable-State Generators. *Entropy* **2014**, *16*, 1396–1413.

[24] Crutchfield, J.P.; Riechers, P.; Ellison, C.J. Exact Complexity: Spectral Decomposition of Intrinsic Computation **submitted**. Santa Fe Institute Working Paper 13-09-028; arXiv:1309.3792 [cond-mat.stat-mech].

[25] *The Entropy of Functions of Finite-state Markov Chains*, Vol. 28, Publishing House of the Czechoslovak Academy of Sciences, Prague, 1957. Held at Liblice near Prague from November 28 to 30, 1956.

[26] Strelioff, C.C.; Crutchfield, J.P. Bayesian Structural Inference for Hidden Processes. *Physical Review E* **2014**, *89*, 042119.

[27] Crutchfield, J.P. Between Order and Chaos. *Nature Physics* **2012**, *8*, 17–24.

[28] Marzen, S.; Crutchfield, J.P. Time Resolution Dependence of Information Measures for Spiking Neurons: Scaling and Universality **2015**. submitted.

[29] Marzen, S.; Crutchfield, J.P. Long-Range Memory in Stationary Renewal Processes **2014**. In preparation.

[30] Ephraim, Y.; Merhav, N. Hidden Markov Processes. *IEEE Trans. Info. Th.* **2002**, *48*, 1518–1569.

[31] Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, second ed.; Wiley-Interscience: New York, 2006.

[32] Crutchfield, J.P.; Young, K. Inferring Statistical Complexity. *Phys. Rev. Let.* **1989**, *63*, 105–108.

[33] Shalizi, C.R.; Crutchfield, J.P. Computational Mechanics: Pattern and Prediction, Structure and Simplicity. *J. Stat. Phys.* **2001**, *104*, 817–879.

[34] Paz, A. *Introduction to Probabilistic Automata*; Academic Press: New York, 1971.

[35] Rabiner, L.R.; Juang, B.H. An Introduction to Hidden Markov Models. *IEEE ASSP Magazine* **1986**, *January*.

[36] Rabiner, L.R. A Tutorial on Hidden Markov Models and Selected Applications. *IEEE Proc.* **1989**, *77*, 257.

[37] Crutchfield, J.P.; Ellison, C.J.; Mahoney, J.R. Time's Barbed Arrow: Irreversibility, Crypticity, and Stored Information. *Phys. Rev. Lett.* **2009**, *103*, 094101.

[38] Ellison, C.J.; Mahoney, J.R.; James, R.G.; Crutchfield, J.P.; Reichardt, J. Information Symmetries in Irreversible Processes. *CHAOS* **2011**, *21*, 037107.

[39] Yeung, R.W. *Information Theory and Network Coding*; Springer: New York, 2008.

[40] James, R.G.; Ellison, C.J.; Crutchfield, J.P. Anatomy of a Bit: Information in a Time Series Observation. *CHAOS* **2011**, *21*, 037109.

[41] Bell, A.J. The Co-information Lattice. Proceedings of the Fifth International Workshop on Independent Component Analysis and Blind Signal Separation; S. Amari, A. Cichocki, S.M.; Murata, N., Eds.; Springer: New York, 2003; Vol. ICA 2003.

[42] Ara, P.M.; James, R.G.; Crutchfield, J.P. The Elusive Present: Hidden Past and Future Correlation and Why We Build Models **2014**. In preparation.

[43] Crutchfield, J.P.; Young, K. Computation at the Onset of Chaos. Entropy, Complexity, and the Physics of Information; Zurek, W., Ed.; Addison-Wesley: Reading, Massachusetts, 1990; Vol. VIII, *SFI Studies in the Sciences of Complexity*, pp. 223 – 269.

[44] Bialek, W.; Nemenman, I.; Tishby, N. Predictability, Complexity, and Learning. *Neural Comp.* **2001**, *13*, 2409–2463.

[45] Crutchfield, J.P. The Calculi of Emergence: Computation, Dynamics, and Induction. *Physica D* **1994**, *75*, 11–54.

[46] Ellison, C.J.; Mahoney, J.R.; Crutchfield, J.P. Prediction, Retrodiction, and the Amount of Information Stored in the Present. *J. Stat. Phys.* **2009**, *136*, 1005–1034.

[47] Godreche, C.; Majumdar, S.N.; Schehr, G. Statistics of the Longest Interval in Renewal Processes. *J. Stat. Mech: Theory and Experiment* **2015**, *March*, P03014.

[48] Birch, J.J. Approximations for the Entropy for Functions of Markov Chains. *Ann. Math. Statist.* **1962**, *33*, 930–938.

[49] Riechers, P.M.; Varn, D.P.; Crutchfield, J.P. Pairwise Correlations in Layered Close-Packed Structures. *Acta Cryst. A* **2015**, *71*, 423–443.

[50] Still, S.; Ellison, C.J.; Crutchfield, J.P. Optimal Causal Inference: Estimating stored information and approximating causal architecture. *Chaos* **2010**, *20*, 037111.

[51] Crutchfield, J.P. Critical Computation, Phase Transitions, and Hierarchical Learning. Towards the Harnessing of Chaos; Yamaguti, M., Ed.; Elsevier Science: Amsterdam, 1994; pp. 29–46.

[52] Crutchfield, J.P.; Feldman, D.P. Statistical Complexity of Simple One-Dimensional Spin Systems. *Phys. Rev. E* **1997**, *55*, R1239–R1243.

[53] Feldman, D.P.; Crutchfield, J.P. Structural Information in Two-Dimensional Patterns: Entropy Convergence and Excess Entropy. *Phys. Rev. E* **2003**, *67*, 051103.

[54] Tchernookov, M.; Nemenman, I. Predictive Information in a Nonequilibrium Critical Model. *J. Stat. Phys.* **2013**,

153, 442–459.

[55] Zipf, G.K. *The Psycho-Biology of Language: An Intro- duction to Dynamic Philology*; Houghton Mifflin Com- pany: Boston, Massachusetts, 1935.

[56] Mandelbrot, B. An informational theory of the statistical structure of languages. In *Communication Theory*; Jack- son, W., Ed.; Butterworths: London, 1953; pp. 486–502.

[57] Miller, G.A. Some effects of intermittent silence. *Amer. J. Psych.* **1957**, *70*, 311–314.

[58] Beggs, J.M.; Plenz, D. Neuronal Avalanches in Neocor- tical Circuits. *J. Nuerosci.* **2003**, *23*, 11167–11177.

[59] Dedeo, S. Collective Phenomena and Non-Finite State Computation in a Human Social System. *PLoS One* **2013**, *8*, e75808.

[60] Cessac, B.; Cofre, R. Spike Train statistics and Gibbs Distributions. *arxiv* **2013**, *1302.5007*.