Controlling epidemic extinction using early warning signals

Walter Ullon[1,2] · Eric Forgoston[1]

Abstract

As the recent COVID-19 pandemic has shown us, there is a critical need to develop new approaches to monitoring the outbreak and spread of infectious disease. Improvements in monitoring will enable a timely implementation of control measures, including vaccine and quarantine, to stem the spread of disease. One such approach involves the use of early warning signals to detect when critical transitions are about to occur. Although the early detection of a stochastic transition is difficult to predict using the generic indicators of early warning signals theory, the changes detected by the indicators do tell us that some type of transition is taking place. This observation will serve as the foundation of the method described in the article. We consider a susceptible–infectious–susceptible epidemic model with reproduction number $R_0 > 1$ so that the deterministic endemic equilibrium is stable. Stochastically, realizations will fluctuate around this equilibrium for a very long time until, as a rare event, the noise will induce a transition from the endemic state to the extinct state. In this article, we describe how metric-based indicators from early warning signals theory can be used to monitor the state of the system. By measuring the autocorrelation, return rate, skewness, and variance of the time series, it is possible to determine when the system is in a weakened state. By applying a control that emulates vaccine/quarantine when the system is in this weakened state, we can cause the disease to go extinct earlier than it otherwise would without control. We also demonstrate that applying a control at the wrong time (when the system is in a non-weakened, highly resilient state) can lead to a longer extinction time than if no control had been applied. This feature underlines the importance of determining the system's state of resilience before attempting to affect its behavior through control measures.

Keywords Epidemic extinction · Early warning signals · Stochastic epidemic model · Control

1 Introduction

Recent events concerning the outbreak of infectious diseases such as Ebola virus disease [1–3], Zika virus [4–6], a particularly deadly strain of Influenza [7,8], and COVID-19 [9] have brought to light the need to develop effective and integrated measures to both monitor and control the spread of epidemics. Only through such an integrated approach will one be able to prevent a regional outbreak from becoming a global pandemic and to mitigate the damage done to the affected areas.

One such approach involves the use of early warning signals (EWS) [10] to detect if a critical transition is approaching. These critical transitions can be found in ecosystems [10], financial systems [11], and the Earth's climate system [12], and may be caused by the system crossing a bifurcation or through stochastic-induced switching [13]. In particular, for problems of disease emergence and ecological tipping points, it has been suggested that a noise-induced switch from the endemic state/carrying capacity to the extinct state can be anticipated if data is collected with sufficient frequency [14,15].

However, this has proven difficult in practice since, for instance, disease emergence is characterized by low prevalence and is often complicated by amplification of transients and oscillatory dynamics [15]. Also, it is widely held that

✉ Eric Forgoston
eric.forgoston@montclair.edu

Walter Ullon
ullonw1@gmail.com

1 Department of Applied Mathematics and Statistics, Montclair State University, 1 Normal Avenue, Montclair, NJ 07043, USA

2 Present Address: EZOPS, Inc., 463 7th Avenue, #1504, New York, NY 10018, USA

noise-induced transitions cannot be detected since there is no change in the shape of the potential function, and no change in the eigenvalue of the mean field model [14,16]. Furthermore, insights from large deviation theory [17,18] cast doubt on the effectiveness, accuracy, and feasibility of anticipating stochastic transitions using EWS [19–21].

Nevertheless, extensive research shows that generic early warning signals exist in a wide variety of systems, and can in fact be the harbingers of critical transitions [10,22]. For instance, the presence of leading indicators such as autocorrelation and variance has been found to increase prior to climatic transitions as well as before regime changes in lake food webs [16,23–25]. And although, as previously stated, the early detection of a stochastic transition is difficult to predict by the sole reliance on the generic indicators developed in the theory of EWS, the changes detected by EWS indicators do tell us that some type of transition is taking place, be it critical or otherwise [26]. This observation will serve as the foundation of this article, as we seek to test and develop methods that introduce control measures at critical stages in the dynamics of the epidemic model we consider.

We consider a susceptible–infectious–susceptible (SIS) epidemic model and employ the metric-based indicators developed in the theory of EWS to monitor the state of the system. By doing this, we are able to ascertain the system's overall resilience to control measures at different points in time. We seek to drive the infectious disease to extinction in an amount of time that is shorter than the noise-induced extinction that would occur in the absence of control. This will be accomplished by perturbing the system at specific points corresponding to states of low resilience, *i.e.,* a weakened system state.

In short, whereas previous attempts focused on the prediction of critical transitions in stochastic systems, we attempt to induce them. We hypothesize that perturbing the system while in a weakened state will precipitate critical transitions, while the same approach applied during periods of increased resilience will be met with resistance, failing to yield similar results. Our hypothesis stems from results that show that many stochastic systems do in fact behave like systems undergoing critical transition due to critical slowing down (CSD) in the mean field model [14].

We test our hypothesis by simulating stochastic data from the SIS epidemic model using a type of Monte Carlo method [27]. This time series data is subjected to EWS analysis [23,24] using metric-based indicators, which quantify changes in statistical properties of the time series. In particular, one can use the autocorrelation, return rate, skewness, and variance to yield actionable information which is then used to introduce quarantine/immunization measures to effect the desired result [23].

Consistent with our proposition, we show that removing a portion of the infected population in the SIS model at a low resilience state through quarantine/vaccine is an effective measure for inducing early extinction events in the epidemic. Mean extinction times (MTE) are computed for the various cases and demonstrate agreement with the hypothesis. Results also warn against the indiscriminate application of controls when the system is in a high resilience state.

2 Early warning signals theory

Early warning signals (EWS) can be thought of as the direct product of the need to understand the phenomenon of critical transitions. In many instances, the critical transition corresponds to a catastrophic bifurcation in which the qualitative behavior of the system changes abruptly in response to changes in one or more external parameters. It commonly arises in deterministic systems with alternative equilibrium states (Fig. 1). However, as demonstrated in this article, the transition can be noise-induced wherein stochasticity causes a transition from a deterministically stable equilibrium to a deterministically unstable equilibrium (e.g., a transition from the endemic state to the extinct state found in epidemic models for a reproduction number, $R_0 > 1$) [1,18,28,29].

In addition, critical transitions often exhibit "threshold behavior," which is accompanied by changes in the system properties that become more extreme as it moves closer to the tipping point. These changes include but are not limited to: increasing autocorrelation, decreased return rate, increasing variance, increasing skewness, and flickering [30]. The gradual change in these collective properties is generally termed critical slowing down (CSD). In essence, it is CSD that allows early warning signals to detect critical transitions in progress. Figure 1 illustrates the role of CSD with respect to the dynamical landscape.

It is worth noting again that there are different types of transitions, some catastrophic, some not, some showing CSD and some not, and some induced by noise. Since we are considering stochastic data, it is important to realize that the stochastic perturbations will influence the EWS statistics, and it is precisely the careful monitoring of these metric-based indicators that will serve as the basis for our analysis. In addition, while CSD is typically associated with undesirable outcomes, such as desertification [31], lake eutrophication [32], and coral reef collapse [33], here we take advantage of the resulting dynamical consequence of CSD to accelerate critical transitions in systems where such outcomes are desirable, such as the extermination of pest populations, and the local extinction of epidemics.
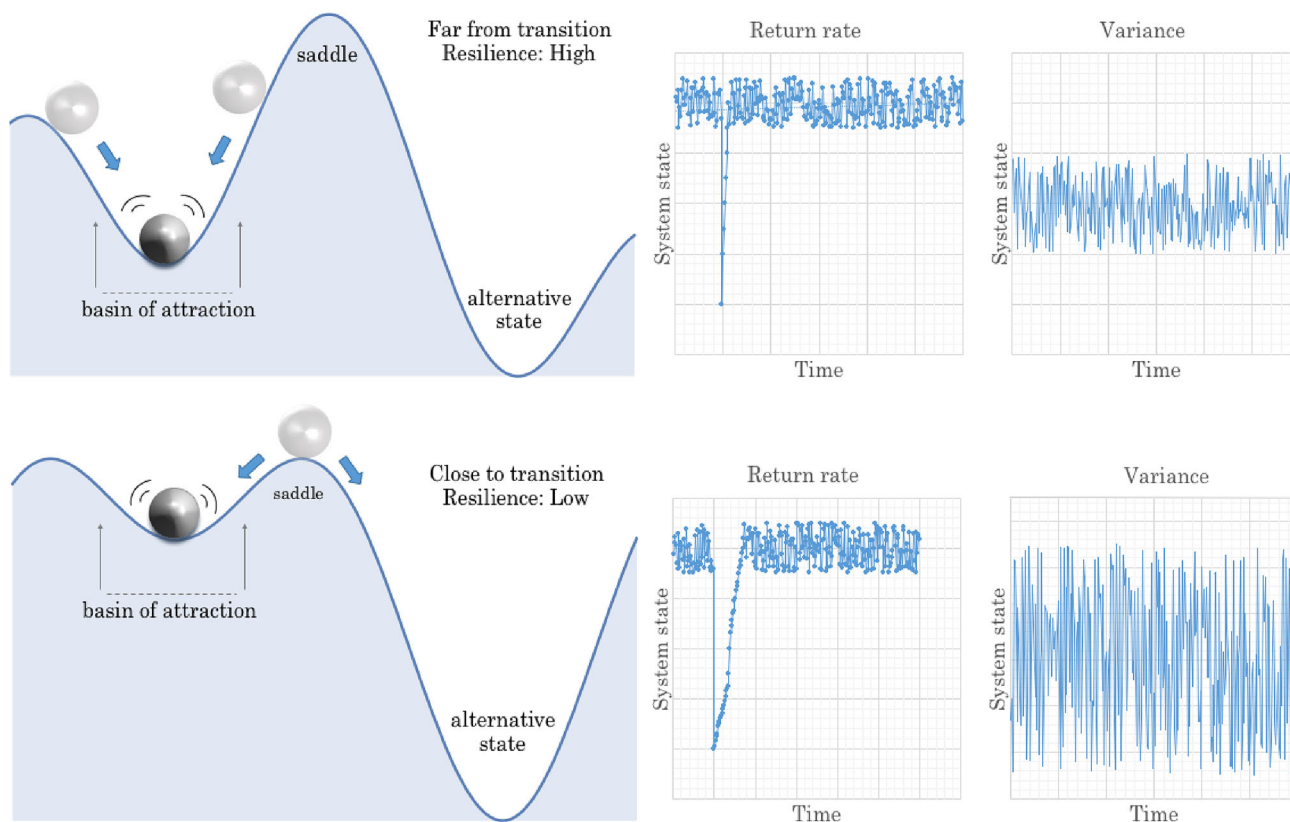
Fig. 1 A sample dynamical landscape for a system with alternative equilibrium states in which the critical transition is brought about by a deterministic bifurcation. Although the *SIS* epidemic model discussed in this article does not possess alternative equilibrium states, we show that the EWS metric indicators can still be used. (Top left) The dynamical landscape far from the transition. A deep basin of attraction makes the switch to the alternative state due to small perturbations difficult. (Top right) Far from the transition, the return rate is fast, and the variance is relatively small. (Bottom left) Flattening of the dynamical landscape due to critical slowing-down can cause even small perturbations to induce a switch to an alternative state. (Bottom right) Close to the transition, the return rate slows down, while the variance increases
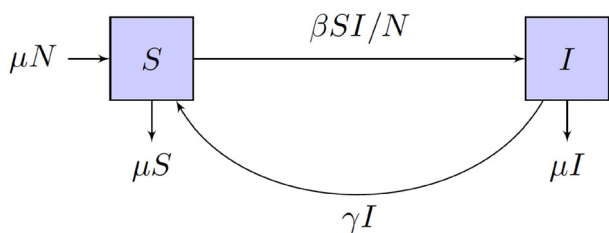
Fig. 2 *SIS* compartmental diagram. Individuals are born susceptible with rate μ, susceptible individuals become infectious with contact rate β, infectious individuals recover and become re-susceptible with rate γ, and both susceptible and infectious individuals die with rate μ

3 *SIS* epidemic model

The *SIS* epidemic model is constructed by dividing the population into two classes of individuals, namely susceptible S and infectious I. Individuals are born into the susceptible population with birth rate μ. A susceptible individual can perish with death rate μ and can become infected with the disease, with contact rate β. An infectious individual recovers and becomes re-susceptible with rate γ, and can perish with death rate μ. The units of all rates are per year. Figure 2 illustrates the SIS compartmental diagram, and the deterministic mean-field equations are given as

$$\frac{dS}{dt} = \frac{-\beta S I}{N} + \mu(N - S) + \gamma I, \tag{1}$$

$$\frac{dI}{dt} = \frac{\beta S I}{N} - \mu I - \gamma I, \tag{2}$$

where N is the constant population size.

Equations (1 and 2) have two steady states. The disease-free or extinct state has no infectious individuals and is given as

$$(S_0, I_0) = (N, 0). \tag{3}$$

The infection is maintained at the endemic disease state, which is given as

$$(S_e, I_e) = \left(\frac{N}{R_0}, N \left(1 - \frac{1}{R_0} \right) \right), \tag{4}$$

Table 1 SIS Transitions and Rates

Event	Transition	S Rate	I Rate
Birth	$S \xrightarrow{\mu} 2S$	μN	
Infection	$2S \xrightarrow{\beta} S, I \xrightarrow{\beta} 2I$	$\beta SI/N$	$\beta SI/N$
Recovery	$S \xrightarrow{\gamma} 2S, 2I \xrightarrow{\gamma} I$	γI	γI
Death	$S, I \xrightarrow{\mu} \varnothing$	μS	μI

Table 2 Parameter values used in the time series simulations of the SIS model

Parameter Name	Type	Value
μ	Birth/death rate	0.02
β	Contact rate	1000
γ	Removal/recovery rate	99.98
R_0	Basic reproductive number	10
N	Initial population size	40
–	No. of iterations	10^7

where $R_0 = \frac{\beta}{\gamma + \mu}$ is the basic reproductive number [34]. The stability of these two fixed points is determined by the value of the reproductive number R_0, which can be thought of as the average number of new infectious individuals that one infectious individual generates over the course of the infectious period in an entirely susceptible population. For $R_0 > 1$, the extinct state is unstable while the endemic state is stable. Note that since the model is deterministic, a population at the attracting endemic state can never go extinct. However, the inclusion of demographic stochasticity will induce a large fluctuation that brings the population into the extinct state.

The corresponding stochastic population model is represented by the transition processes of birth, death, infection, and recovery [18,35,36]. The associated rates are given in Table 1.

4 Methods

In this section, we describe the methodology used to generate the data as well as the statistical methods used to understand the application of the control at different resiliencies. Specifically, Sect. 4.1 contains details about the Monte Carlo method that is used to generate 1000 stochastic realizations (time series) of the SIS epidemic model, while Sect. 4.2 describes the metric-based indicators that are used to capture changes in the properties of the time series. Based on the resulting autocorrelation data for each time series, we determine the low and high thresholds that correspond, respectively, to regions of high and low resilience (Sect. 4.3). Lastly, Sect. 4.4 describes the control mechanism and the determination of mean extinction times.

4.1 Simulation of the stochastic *SIS* model

To generate a solution of this stochastic model, where the demographic noise is internal to the system, we use the Doob–Gillespie algorithm (also known as the Gillespie algorithm or Gillespie's stochastic simulation algorithm (SSA)) [27,37]. The algorithm is a type of Monte Carlo method that was originally proposed by Kendall [38] for simulating birth–death processes and was popularized by Gillespie [27] as a useful method for simulating chemical reactions based on

molecular collisions. The results of a Gillespie simulation is a stochastic trajectory that represents an exact sample from the probability function that solves the master equation. Therefore, the method can be used to simulate population dynamics where molecular collisions are replaced by individual events and interactions including birth, death, and infection [18].

Let $\mathbf{x} = (x_1, \ldots, x_n)^T$ denote the state variables of a system, where x_j provides the number of individuals in state x_j at time t. The first step of the algorithm is to initialize the number of individuals in the population compartments $\mathbf{x}_0$. For a given state $\mathbf{x}$ of the system, one calculates the transition rates (birth rate, death rate, contact rate, etc.) denoted as $a_i(\mathbf{x})$ for $i = 1 \ldots l$, where l is the number of transitions. Thus, the sum of all transition rates is given by $a_0 = \sum_{i=1}^{l} a_i(\mathbf{x})$.

Random numbers are generated to determine both the next event to occur as well as the time at which the next event will occur. One simulates the time τ until the next transition by drawing from an exponential distribution with mean $1/a_0$. This is equivalent to drawing a random number r_1 uniformly on $(0, 1)$ and computing $\tau = (1/a_0) \ln (1/r_1)$. During each random time step, exactly one event occurs. The probability of any particular event taking place is equal to its own transition rate divided by the sum of all transition rates $a_i(\mathbf{x})/a_0$. A second random number r_2 is drawn uniformly on $(0, 1)$, and it is used to determine the transition event that occurs. If $0 < r_2 < a_1(\mathbf{x})/a_0$, then the first transition occurs; if $a_1(\mathbf{x})/a_0 < r_2 < (a_1(\mathbf{x}) + a_2(\mathbf{x}))/a_0$, then the second transition occurs, and so on.

Lastly, both the time step and the number of individuals in each compartment are updated, and the process is iterated until the disease goes extinct or until the simulation time has been exceeded [18]. A sample realization of the stochastic SIS model using the transitions in Table 1 and the parameter values shown in Table 2 is shown in the top panel of Fig. 3. The Table 2 parameter values are used to generate all the simulated data used throughout this article, and the code used to generate this realization as well as all the simulated data can be found in the github repository https://github.com/Walter-Ullon/Stochastic-SIS.

Fig. 3 (Top) Stochastic time series for the SIS model generated using the Gillespie algorithm as described in the text. (Middle) Autocorrelation, and (Bottom) return rate for the stochastic time series shown in the top panel computed using the Early Warning Signals Toolbox [39]. The green band indicates the ideal time to apply the control (low resilience), while the red band indicates where the control could lead to undesired results (high resilience). These, and other, statistical metric-based indicators provide an accurate picture of the state of the SIS system
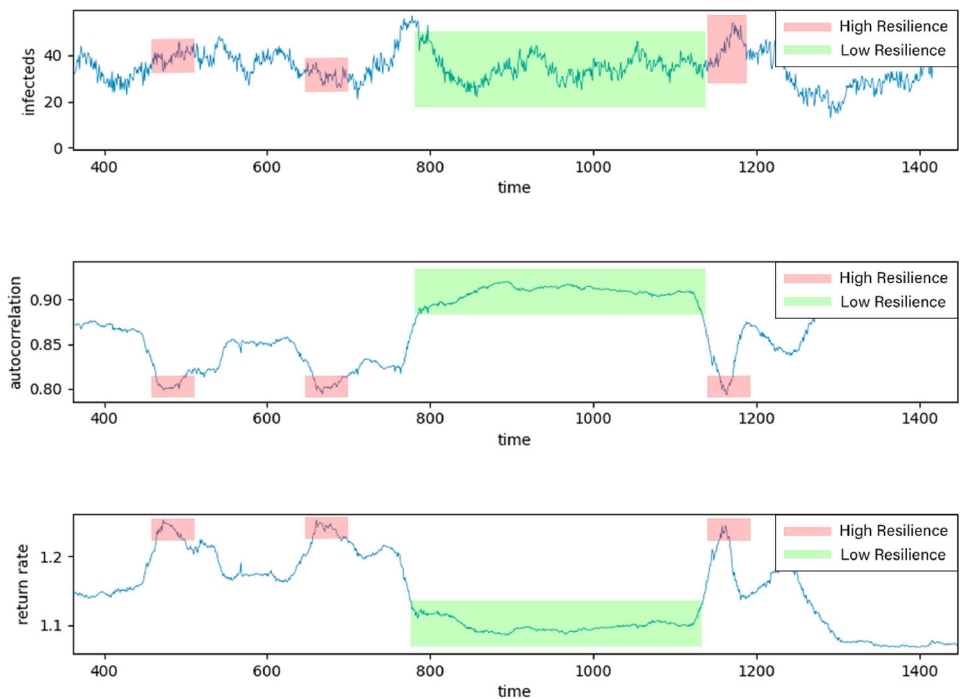
4.2 Statistical analysis: metric-based indicators

Statistical analysis of the simulated stochastic time series data is performed using the Early Warning Signals Toolbox for R [39]. As stated by its developers, the toolbox provides methods for estimating statistical changes in time series data that can be used for identifying nearby critical transitions [23].

Metric-based indicators capture changes in the properties of an observed time series of a system and quantify changes in the statistical properties of the time series without attempting to fit the data with a specific model structure [23]. These methods include a variety of statistical measures including autocorrelation and spectral properties, variance, skewness and kurtosis, detrended fluctuation analysis, and conditional heteroskedasticity.

While the Early Warning Signals Toolbox also contains robust model-based methods to analyze time series data, in our work we seek to rely only on metric-based methods as this provides us with more flexibility in the application of the theory across various dynamical systems. From the aforementioned indicators, autocorrelation and return rate are the simplest methods to measure critical slowing down. An increase in autocorrelation indicates that the state of the system has become increasingly similar between consecutive observations, while a decrease in the return-rate signals the system's inability to return to its previous state. Both measures provide evidence of critical slowing down [23].

The Early Warning Signals Toolbox measures lag-1 autocorrelation in three separate ways, namely the autocorrelation function, an autoregressive model of order 1, and return rate

[23]. In this work, we have relied on autocorrelation found using the autoregressive model using a rolling window of 25% of the time series based on the idea that indicators should be estimated as data becomes available. The autocorrelation found using the autocorrelation function is very similar for the time series generated in this work.

Figure 3 shows two statistical indicators, autocorrelation and return rate, for the associated stochastic SIS time series. As we can observe in Fig. 3, it is possible for these indicators to signal a critical transition in progress, while the time series of SIS infectives still does not switch from the endemic state to the alternative extinct state, i.e., a false positive. However, by controlling the system (i.e., removing a portion of the infected individuals) at the right time, one can precipitate the transition and achieve an early extinction in which extinction occurs earlier than it would have with no control. This loss of resilience is a trademark of CSD. As we shall see, the results obtained by relying on metric-based EWS analysis are both encouraging and sobering.

4.3 Statistical analysis: resilience thresholds

Since our methods call for the introduction of control measures when the system is displaying threshold values in autocorrelation (high/low), it is essential that we define specifically what these thresholds are and the manner in which they were obtained. We note that at this point, we take autocorrelation as our leading metric-based indicator in the analysis and simulations that follow.

Table 3 Threshold values for autocorrelation for the SIS model

Threshold Type	Threshold Value
Maximum	0.90713
Minimum	0.71253
Standard deviation	0.0335
High threshold	0.87366
Low threshold	0.74601

To ensure robust statistical estimates, one thousand stochastic realizations of the SIS model were computed using the Gillespie algorithm outlined in Sect. 4.1. The autocorrelation data was computed for each of the thousand realizations in the manner described in Sect. 4.2, and the autocorrelation time series was averaged to obtain a mean autocorrelation value specific to the time series being considered. The end result of this process is one thousand autocorrelation means, which were in turn averaged to obtain a clearer idea of the distribution and location of the true mean autocorrelation for the SIS model.

This process was carried out in the same manner for the maximum and minimum autocorrelation values of each time series, which yielded one thousand maximum/minimum data points, which were subsequently averaged. Ultimately, these measures provided us with a statistical approximation for the location of the high and low autocorrelation thresholds.

To account for the inherent variation present in these simulations, we also obtained the mean standard deviation for the model which was used to create "soft thresholds." Thus, as we will see in the next section, when the control was applied at the high/low autocorrelation point, it was done only when these values surpassed the maximum/minimum $\mp$ one standard deviation. The choice to employ the soft thresholds was done in order to capture more time series that exhibit threshold behavior, as some realizations never reach or surpass the maximum/minimum thresholds.

The thresholding results are summarized in Table 3 and Fig. 4. These threshold values provide a metric by which to determine how far we are willing to let the system evolve before the application of control measures. In short, we introduce vaccination/quarantine measures in the SIS model whenever the autocorrelation reaches or surpasses the aforementioned threshold values for high/low thresholds.

In general, for a new epidemic outbreak, one does not have vast troves of data available to determine the thresholds. However, it is possible to calculate the threshold value from the beginning of a single stochastic realization corresponding to a specific disease outbreak. One could also find the threshold values using data from previous or similar disease outbreaks.

4.4 Control mechanism: pulsing the population

From our previous discussion of EWS theory, we know that the time series statistics when far from the transition are very different from the statistics when close to the transition. This is reflected by the rise in autocorrelation, variance, and skewness, and inversely by a drop in the return rate. It is precisely at this point where we introduce control measures such as quarantine or vaccines in the SIS model.

The control can be thought of as either an indefinite quarantine or a mass vaccination event. One could also envision a scenario in which quarantined individuals recover and are given a vaccine after quarantine to prevent re-susceptibility. Either way, the "removed" individuals do not interact with the rest of the population (susceptibles or infectious) and do not contribute to the SI dynamics. The control causes a decrease in both the number of infectious individuals and the total population N since $S + I = N$.

Figure 5 demonstrates the effect of control when it is applied at the low and high resilience states. The top panel plot shows the result of a typical run of the stochastic SIS model unaffected by controls or outside measures. Even at time $t = 1200$, the infected population is still fluctuating around the endemic steady state.

Following with our analysis of the metric-based indicators for critical transitions, for the middle plot we realize a second run of the top plot using the same random seed and parameters, where we monitor for high autocorrelation values corresponding to low resilience. Once this threshold is reached (as per the values outlined in Fig. 4), we remove a portion of the population, i.e., the control is applied. The infected population struggles to return to the deterministically stable endemic state and eventually goes extinct. Thus, applying the control at the low-resilience state corresponding to high autocorrelation, variance, skewness, and decrease in return rate suffices to induce a critical transition to the extinct state.

For the bottom plot, we again recreate an exact re-simulation of the stochastic SIS, but this time we monitor the metric indicators for signals corresponding to high resilience, namely low autocorrelation values. Here, the infected population behavior is much different; it is able to sustain the removal of individuals, effectively returning to fluctuate about the endemic state as it did before the control measure. Hence, consistent with our hypothesis, quarantining individuals at the high-resilience state (low autocorrelation, high return rate) fails to precipitate an early extinction event.

In addition to the individual realizations presented in Fig. 5, we performed the same analysis for one thousand realizations of the model utilizing the thresholds outlined in Sect. 4.3. By comparing the mean time to extinction (MTE) of the SIS realizations with and without control measures as well as with a blind test where the control was applied

Fig. 4 Autocorrelation analysis for one thousand stochastic realizations of the SIS model. Note that the "high threshold" value in the legend represents the maximum autocorrelation value minus one standard deviation, while the "low threshold" represents the minimum autocorrelation value plus one standard deviation
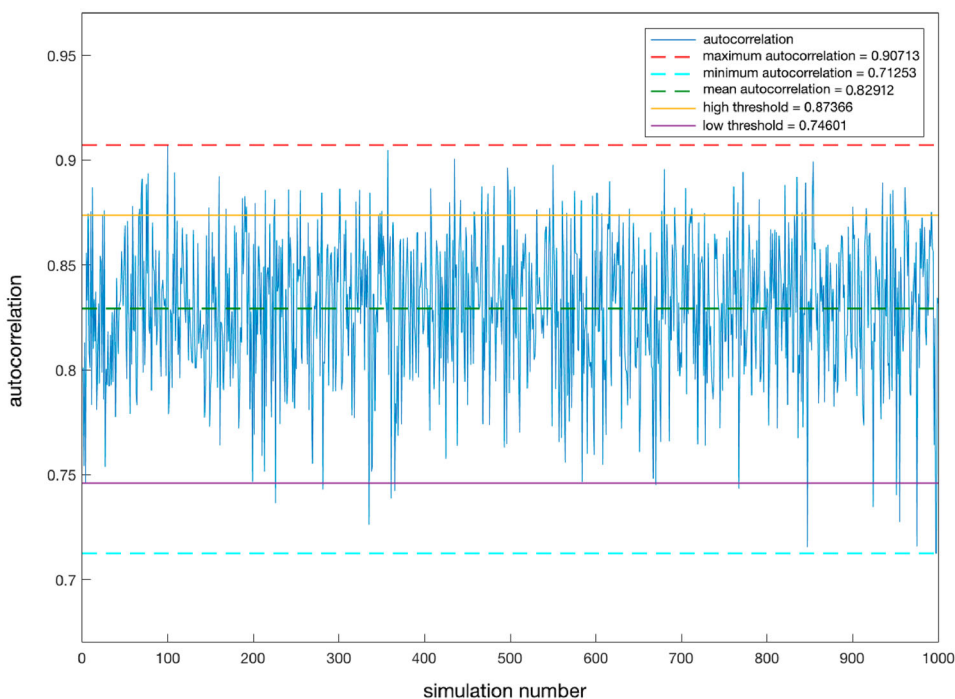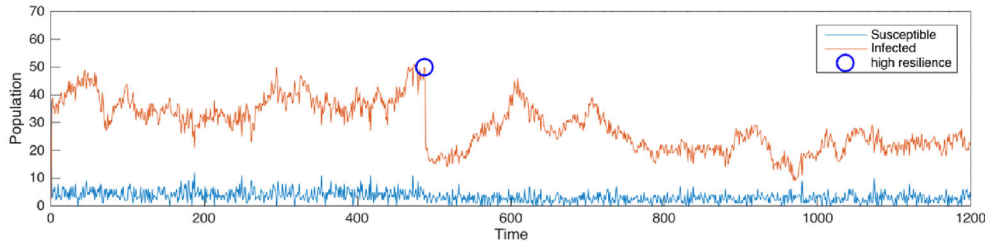

Initial Population Size = 40
Unperturbed
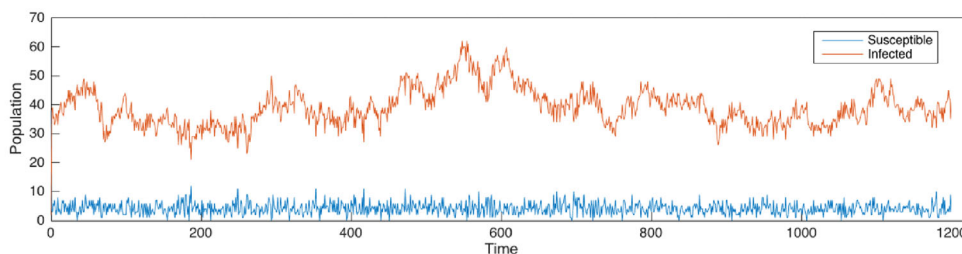
t = 623
Pulse = -29
Population at pulse = 46
Resilience = low
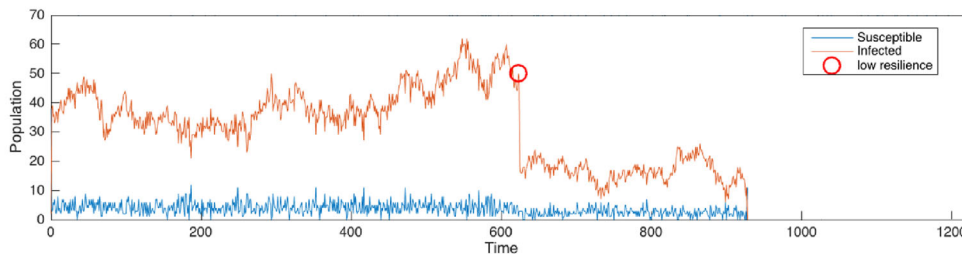
t = 487
Pulse = -29
Population at pulse = 45
Resilience = high

Fig. 5 Time series for the SIS model. (Top) an unperturbed realization of the model. After 10^7 iterations, the epidemic continues to fluctuate about the endemic state (red). (Middle) the control is applied at the low-resilience state (29 of 46 infected individuals are removed), and extinction follows shortly after. (Bottom) the control is applied at the high-resilience state (29 of 45 infected individuals are removed); the system is able to resist the control, recover, and continue to fluctuate about the endemic state
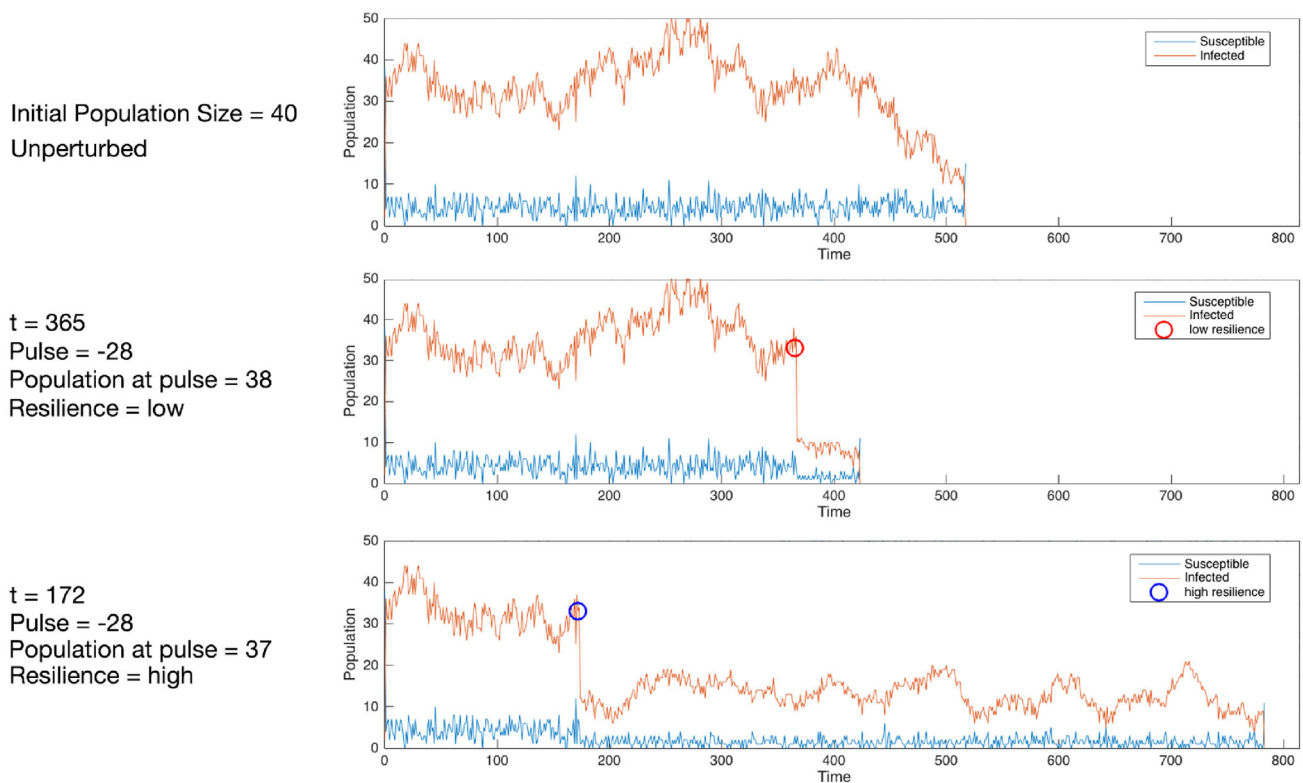
Fig. 6 Time series for the SIS model. (Top) An unperturbed realization of the model where the disease undergoes extinction on its own. (Middle) The control is applied at the low-resilience state (28 of 38 infected individuals are removed), and extinction follows shortly after. (Bottom) The control is applied at the high-resilience state (28 of 37 infected individuals are removed). As in the previous case, the system is able to resist the control and recover. However, it eventually goes extinct at a much later time than if no control measure was applied (see Top panel)

randomly, it is possible to compute a quantitative measure of the effectiveness of the EWS analysis.

Because not all of the individual SIS realizations reached extinction in the prescribed computational bounds (defined by us to be 10^7 random time steps), we separately compare "effective" versus "absolute" MTE. The effective MTE is calculated by averaging the individual extinction times only for the realizations that went extinct after the control was applied and before the maximum number of time steps was reached, whereas the absolute MTE is calculated by incorporating into the average those realizations that had not yet gone extinct despite the application of control (using the maximum time as an "extinction" time). These two types of MTE were computed for four possible scenarios: (i) no control is applied; (ii) the control is applied at the high resilience point; (iii) the control is applied at the low resilience point; and (iv) the control is applied at random. The results are provided in the following section (Sect. 5).

For the purpose of our simulation, we opted to vaccinate/quarantine 40% of the infected population once the thresholds were met. While several removal factors were implemented (30% up to 75%), we note that in our simulations, 40% provided the best results, as a lower fraction

proves ineffective in driving the critical transition, while higher percentages move the system so far away from the endemic state that the system loses the capability to recover and a very rapid extinction occurs—this scenario does little to provide us with a clear understanding of the underlying dynamics and resiliency of the system. In essence, if the control is too small the effect cannot contribute to the governing dynamics in overriding the current state, while if the control is too large, the contribution overwhelms the governing dynamics without yielding any useful information.

5 Results and discussion

As we have seen in Fig. 5, the results of our experiment are encouraging since we have shown that monitoring the system and attacking it with conventional control measures can lead to a significant reduction in the active cycle of the epidemic. Nevertheless, this approach needs to be employed with care.

For instance, the results presented in Fig. 6 paint a perplexing picture. Consistent with the case examined previously, the quarantining/immunization of infected individuals during weakened system states drives it to early extinction. However,

Table 4 Effective mean extinction times for 1,000 realizations (years)

Threshold type	no. times threshold was met	no. times extinction was reached	MTE (effective)
low	770	134	1265.8
high	894	235	1228.6
blind	n/a (random)	153	1237
no control	n/a	1,000	1,819,997.8

application of the control at the high-resilience state exacerbates adverse conditions, ultimately causing an extension of the life cycle of the epidemic. This is at the very least counter intuitive, since reason dictates that when faced with the outbreak of some epidemic the best alternative is to act rather than sit idly by. But as this result shows, doing the right thing at the wrong time can lead to dire consequences.

SIS mean extinction time

After applying the controls to the indicated fraction of the population, we compared the results of applying the control at the high threshold, applying the control at the low threshold, and applying the control randomly (blind test) to the case of no control. From one thousand realizations, we were able to meet the high threshold 894 times, achieving early extinction 235 times, for an MTE of 1228.6 years.

On the other hand, the low threshold was met 770 times, reaching extinction only 134 times for an MTE of 1265.8 years. The results confirm our hypothesis, and show that applying the control at the high threshold is at least 8.9% more effective than the low threshold, resulting in a difference of approximately 37 years in terms of MTE.

As for the blind test, we applied the control at random. This resulted in 153 extinctions for an effective MTE of 1237 years. As expected, the use of random control falls between the results of applying the control at low and high threshold. We expected to perform better than applying the control at the low resilience point. In this case, the MTE is shorter by approximately eight years. We also expected to do worse by applying the control at the high resilience point. Application of control at high resilience leads to a longer MTE by approximately 28 years.

Note that regardless of the difference in MTE between the different scenarios, more extinctions were ultimately achieved by applying the control at the low resilience state versus applying the control at either the high resilience state or at random.

The numbers are far more dramatic when comparing the results of the high threshold control versus no controls measures. To obtain the "no control" MTE, we computed the simulations for 10^{11} time steps. The results are outlined in Table 4.

Table 5 Absolute mean extinction times for 1,000 realizations (years)

Threshold type	MTE (absolute)
low	1,576,282
high	1,392,587
blind	1,541,727
no control	1,819,997.8

Mean extinction time: effective versus absolute

We note that the values outlined in Table 4 represent the effective MTE, which were obtained by aggregating and averaging the final extinction times for only those trials that resulted in extinction post-control. Because those simulations were only computed for 10^7 time steps, some realizations were still fluctuating around the endemic state at the time the simulation ended. Taking those into consideration and using the "no control" figures as a guide, we can easily calculate the absolute MTE. The results are outlined in Table 5, and are consistent with our hypothesis and the theory of EWS.

The results for the *SIS* model show that it is in our best interest to extract information from the system before we attempt to introduce controls, as we could ultimately be extending the same conditions that we are trying to eliminate. Furthermore, extensive simulations show that there is an advantage in introducing controls at the appropriate moment when trying to induce early extinctions. Last, the difference in MTE of the control versus no control cases show us that EWS analysis can be an effective tool in the fight against disease epidemics.

6 Conclusions and future work

Our results show that metric-based EWS analysis is able to capture critical changes in the system state under different conditions, calling to our attention the effectiveness of early warning signals theory to monitor the resilience of dynamic systems across a variety of scenarios.

In addition, the control methods explored, aided by the theory, provide us with the ability to explore novel ways to solve problems of great importance, where the focus is to move away from the prediction of transitions toward the control of outcomes. As we have seen, the results obtained

can be of great value and hopefully they will serve as the launching platform for further studies in this area.

In the future, we seek to extend our work to include the study of additional dynamical models, thus gaining a deeper understanding of the theory of early warning signals and its applicability to real world problems. Also, further scaling of the models presented here to include larger populations, as well as large-scale automated testing would enhance the robustness of the results achieved thus far. This is important as unfortunately, the simulations are computationally expensive and at the moment long runs with large populations are prohibitive. These improvements would enable the application of this work to novel epidemic outbreaks.

Data Availability The code used to generate all the simulated data can be found in the github repository https://github.com/Walter-Ullon/Stochastic-SIS.

Declarations

Conflict of Interest The authors have no conflicts to disclose.

References

1. Nieddu GT, Billings L, Kaufman JH et al (2017) Extinction pathways and outbreak vulnerability in a stochastic Ebola model. J R Soc Interface 14(127):20160,847

2. Timothy JW, Hall Y, Akoi-Boré J et al (2019) Early transmission and case fatality of Ebola virus at the index site of the 2013–16 west African Ebola outbreak: a cross-sectional seroprevalence survey. Lancet Infect Dis 19(4):429–438

3. Tariq A, Roosa K, Mizumoto K et al (2019) Assessing reporting delays and the effective reproduction number: the 2018–19 Ebola epidemic in DRC, May 2018-January 2019. Epidemics 26:128–133

4. Duffy MR, Chen TH, Hancock WT et al (2009) Zika virus outbreak on Yap Island, federated states of Micronesia. N Engl J Med 360(24):2536–2543

5. Musso D, Nilles E, Cao-Lormeau VM (2014) Rapid spread of emerging Zika virus in the Pacific area. Clin Microbiol Infect 20(10):O595–O596

6. Gao D, Lou Y, He D et al (2016) Prevention and control of Zika as a mosquito-borne and sexually transmitted disease: a mathematical modeling analysis. Sci Rep 6(28):070

7. Coburn BJ, Wagner BG, Blower S (2009) Modeling influenza epidemics and pandemics: insights into the future of swine flu (H1N1). BMC Med 7(1):30

8. Chowell G, Echevarría-Zuno S, Viboud C et al (2011) Characterizing the epidemiology of the 2009 influenza A/H1N1 pandemic in Mexico. PLoS Med 8(5):e1000

9. Velavan TP, Meyer CG (2020) The COVID-19 epidemic. Tropical Med Int Health 25(3):278

10. Scheffer M, Bascompte J, Brock WA et al (2009) Early-warning signals for critical transitions. Nature 461(7260):53

11. May RM, Levin SA, Sugihara G (2008) Complex systems: ecology for bankers. Nature 451(7181):893

12. Lenton TM, Held H, Kriegler E et al (2008) Tipping elements in the earth's climate system. Proc Natl Acad Sci 105(6):1786–1793

13. Dutta PS, Sharma Y, Abbott KC (2018) Robustness of early warning signals for catastrophic and non-catastrophic transitions. Oikos 127(9):1251–1263

14. Drake JM (2013) Early warning signals of stochastic switching. Proc R Soc Lond B Biol Sci 280(1766):20130,686

15. O'Regan SM, Drake JM (2013) Theory of early warning signals of disease emergence and leading indicators of elimination. Thyroid Res 6(3):333–357

16. Kuehn C (2011) A mathematical framework for critical transitions: bifurcations, fast-slow systems and stochastic dynamics. Physica D 240(12):1020–1035

17. Freidlin MI, Wentzell AD (1984) Random perturbations of dynamical systems. Springer, New York

18. Forgoston E, Moore RO (2018) A primer on noise-induced transitions in applied dynamical systems. SIAM Rev 60(4):969–1009

19. Boettiger C, Hastings A (2013) No early warning signals for stochastic transitions: insights from large deviation theory. Proc R Soc Lond B Biol Sci 280(1766):20131

20. Boettiger C, Hastings A (2012) Quantifying limits to detection of early warning for critical transitions. J R Soc Interface 9(75):2527–2539

21. Boettiger C, Hastings A (2012) Early warning signals and the prosecutor's fallacy. Proc R Soc Lond B Biol Sci 279(1748):4734–4739

22. Scheffer M (2009) Critical transitions in nature and society. Princeton University Press, Princeton

23. Dakos V, Carpenter SR, Brock WA et al (2012) Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. PLoS One 7(7):e41,010

24. Dakos V, Scheffer M, van Nes EH et al (2008) Slowing down as an early warning signal for abrupt climate change. Proc Natl Acad Sci 105(38):14,308-14,312

25. Livina VN, Lenton TM (2007) A modified method for detecting incipient bifurcations in a dynamical system. Geophys Res Lett 34:L03712

26. Kéfi S, Dakos V, Scheffer M et al (2013) Early warning signals also precede non-catastrophic transitions. Oikos 122(5):641–648

27. Gillespie DT (1977) Exact stochastic simulation of coupled chemical reactions. J Phys Chem 81(25):2340–2361

28. Bauver M, Forgoston E, Billings L (2016) Computing the optimal path in stochastic dynamical systems. Chaos Interdisc J Nonlinear Sci 26(8):083–101

29. Billings L, Forgoston E (2018) Seasonal forcing in stochastic epidemiology models. Ricerche Mat 67(1):27–47

30. Dai L, Vorselen D, Korolev KS et al (2012) Generic indicators for loss of resilience before a tipping point leading to population collapse. Science 336(6085):1175–1177

31. Rietkerk M, Dekker SC, De Ruiter PC et al (2004) Self-organized patchiness and catastrophic shifts in ecosystems. Science 305(5692):1926–1929

32. Wang R, Dearing JA, Langdon PG et al (2012) Flickering gives early warning signals of a critical transition to a eutrophic lake state. Nature 492(7429):419

33. Knowlton N (1992) Thresholds and multiple stable states in coral reef community dynamics. Am Zool 32(6):674–682

34. Brauer F, Castillo-Chavez C, Castillo-Chavez C (2012) Mathematical models in population biology and epidemiology, vol 1. Springer, New York

35. Doering CR, Sargsyan KV, Sander LM (2005) Extinction times for birth-death processes: exact results, continuum asymptotics, and

the failure of the Fokker-Planck approximation. Multiscale Model Simul 3(2):283–299

36. Dykman MI, Schwartz IB, Landsman AS (2008) Disease extinction in the presence of random vaccination. Phys Rev Lett 101(7):078,101

37. Doob JL (1942) Topics in the theory of Markoff chains. Trans Am Math Soc 52(1):37–64

38. Kendall DG (1950) An artificial realization of a simple birth-and-death process. J R Stat Soc Ser B (Methodol) 12(1):116–119

39. Dakos V (2019) Early warning signals toolbox: a user's for detecting critical transitions in time series and spatial data http://www.early-warning-signals.org/, [accessed online 03-August-2021]