



Perspective

Monitoring the Path to the Elimination of Infectious Diseases

John M. Drake^{1,2,3,*}  and Simon I. Hay^{4,5}

¹ Odum School of Ecology, University of Georgia, Athens, GA 30602-2202, USA

² Center for the Ecology of Infectious Diseases, University of Georgia, Athens, GA 30602-2202, USA

³ Department of Zoology, University of Oxford, Oxford OX2, UK

⁴ Institute for Health Metrics and Evaluation, University of Washington, Seattle, WA 98121, USA; sihay@uw.edu

⁵ Big Data Institute, Li Ka Shing Centre for Health Information and Discovery, University of Oxford, Oxford OX3 7LF, UK

* Correspondence: jdrake@uga.edu; Tel.: +1-706-583-5539

Received: 2 March 2017; Accepted: 21 June 2017; Published: 26 June 2017

Abstract: During the endgame of elimination programs, parasite populations may exhibit dynamical phenomena not typical of endemic disease. Particularly, monitoring programs for tracking infection prevalence may be hampered by overall rarity, the sporadic and unpredictable timing and location of outbreaks, and under-reporting. A particularly important problem for monitoring is determining the distance that must be covered to achieve the elimination threshold at an effective reproduction number less than one. In this perspective, we suggest that this problem may be overcome by measuring critical slowing down. Critical slowing down is a phenomenon exhibited by nonlinear dynamical systems in the vicinity of a critical threshold. In infectious disease dynamics, critical slowing down is expressed as an increase in the coefficient of variation and other properties of the fluctuations in the number of cases. In simulations, we show the coefficient of variation to be insensitive to under-reporting error and therefore a robust measurement of the approach to elimination. Additionally, we show that there is an inevitable delay between the time at which the effective reproduction number is reduced to below one and complete elimination is achieved. We urge that monitoring programs include dynamical properties such as critical slowing down in their metrics for measuring achievement and avoid withdrawing control activities prematurely.

Keywords: bifurcation delay; critical slowing down; elimination; endgame; smallpox

1. Introduction

Eradication of an infectious disease is the ultimate success in public health [1]. Smallpox is the only infectious disease of humans to have been eradicated to date [2]. Significant, albeit fragile, progress has been made toward the eradication of polio [3], dracunculiasis [4], and yaws [5]. Moreover, local elimination plans are ongoing for malaria [6], lymphatic filariasis [7,8], measles [9], rubella [10], onchocerciasis [11,12], schistosomiasis [13], and trypanosomiasis [14].

Epidemiology, surveillance, and the changing effectiveness of control actions all present unique “end game” challenges in eradication and elimination efforts [1,15]. A particularly intractable problem has been maintaining the political will to continue interventions as the number of new cases declines, partly because of the declining marginal returns on investment when measured in terms of the reduction in cases and partly because of the difficulty of monitoring sparsely occurring disease episodes among largely inaccessible populations [14,16,17]. If surveillance systems for these diseases could identify the approach to ultimate elimination, in contrast to merely marginal declines in the number of new cases, philanthropic funders, public health officials, and others might be better armed

to both justify and target efforts during the “longest last mile,” especially if it can be shown that the tipping point has been crossed and all that remains is to sustain the gains already achieved and allow transmission to run its course [18–20].

Here we propose a new approach to monitoring the path to elimination, based on recent developments in the theory of leading indicators for dynamical transitions [21]. Leading indicators are statistical patterns exhibited by certain complex systems as they approach a tipping point between alternative states, primarily due to *critical slowing down*, in which the tendency for a dynamical system to return to its steady state is increasingly weakened [21]. In epidemiology, endemism and disease-free status are alternative states, with the vaccination threshold acting as the tipping point between [22]. Although not all switching among alternative states is preceded by leading indicators [23], the transcritical bifurcation that occurs in disease elimination does, implying that the sequence of case reports will contain statistical signatures of the approach to elimination [22,24]. Thus, leading indicators based on critical slowing down may provide a means to measure how successful elimination efforts have been, particularly by indicating if the tipping point has already been crossed or is close at hand [24].

Distinguishing these conditions is clearly important for policy. If the threshold has already been crossed, elimination only requires that the current policy be maintained while transient chains of transmission are allowed to peter out; alternatively, if the threshold is near but not yet crossed, then increases in effort or pulse interventions such as supplementary vaccination campaigns may be required to push the contagious process over the tipping point. Our previous studies developed the theory of critical slowing down in disease transmission systems [22,24], but did not address practical issues such as under-reporting or extrapolating from measured signals to predict the time of threshold crossing.

Smallpox illustrates the potential of this theory for monitoring progress toward disease elimination. Smallpox is an acute and often lethal infection, and was globally eradicated in the wild in 1977 [2]. Nation states differed in both the delivery of vaccination and the speed of elimination, but it was frequently on the time scale of decades, with a rapid decline followed by a long tail. Presumably, the vaccination threshold was crossed sometime before or during the decline phase. The simple question is, when?

2. Model

We illustrate critical slowing down in smallpox elimination with a simple model of transmission and vaccination (Figure 1). This model begins in the endemic state and evolves stochastically assuming a policy of newborn vaccination beginning in year 350. Vaccination programs are imperfect and require time to ramp up. To realistically represent the roll-out of a vaccination campaign, we assumed that disease prevalence and immunity at the start of the simulation was at the natural endemic equilibrium. Vaccine coverage started at zero and increased asymptotically toward 96% (Figure 1). The time the vaccination threshold was crossed is plotted as a vertical dashed line. As expected, before vaccination begins, the dynamics fluctuate around an endemic equilibrium. During the campaign, the number of cases declined, crossing the deterministic vaccination threshold in year 381. The pathogen was not eliminated until much later. This long tail is an example of *bifurcation delay*.

Both the approach to the tipping point (dashed line in Figure 1) and the transmission tail are of interest. From a health policy standpoint, the distinction between these two—and the ability to recognize which phase one is in—is key. If vaccination coverage can be maintained long enough once the threshold has been reached, then elimination is sure to occur eventually. By contrast, if the threshold has not been reached, then elimination is very unlikely and any stochastic extinction that may occur is in a fragile state that can be reversed at any time. During the time between crossing the threshold and elimination, it is important not to reduce vaccination pressure, since ceasing vaccination during this period would allow the disease to return to pre-threshold levels and the gains of the campaign will have been wasted despite having achieved the necessary coverage to ensure elimination.

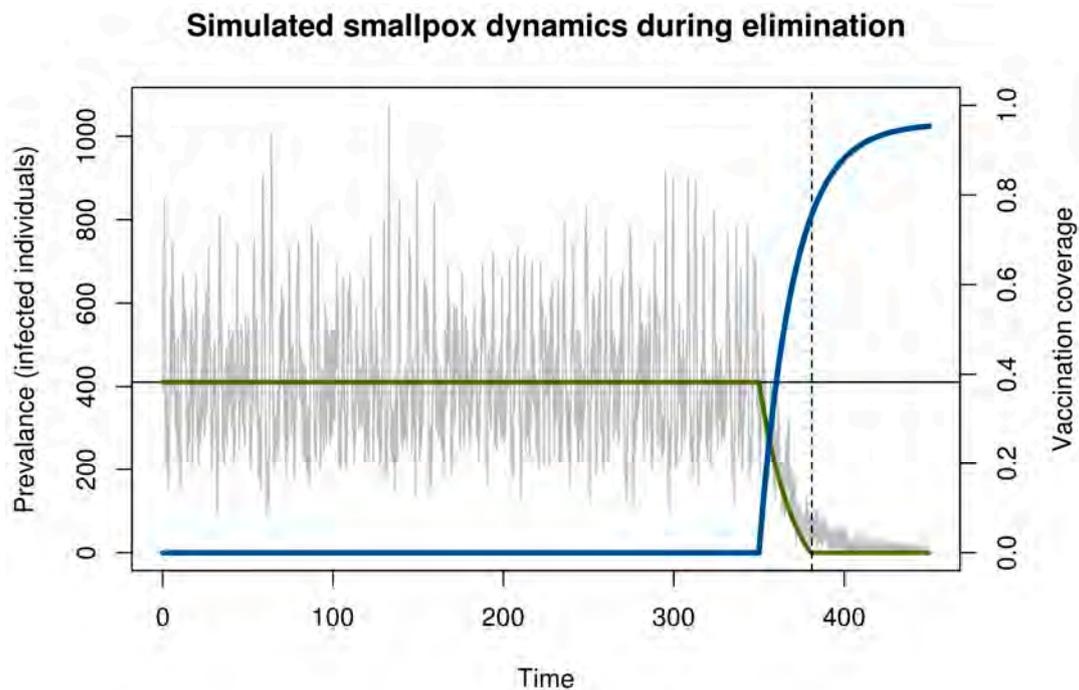


Figure 1. Simulation of smallpox elimination through vaccination. The gray line shows the number of infected individuals over time. The blue line shows vaccination coverage, which approaches a maximum at 0.96. The green line shows the deterministic equilibrium that would be achieved if the vaccination rate was held constant and the time-dependent value. Disease dynamics are given by a stochastic SIR model with states S , I , and R for the number of susceptible, infected, and immune persons, respectively; mean field equations $dS/dt = \mu(S + I + R)(1 - \rho) - \beta SI / (S + I + R) - \xi S - \mu S$, $dI/dt = \beta SI / (S + I + R) + \xi S - \mu I - \gamma I$, and $dR/dt = \gamma I + \mu(S + I + R)(\rho) - \mu R$; and parameters for transmission $\beta = R_0(\gamma + \mu) \approx 121.7$, recovery $\gamma = 365/12 \approx 30.4$, demography $\mu = 1/60$, externally acquired infection $\xi = 0.001$, speed of vaccination roll out $\sigma = -0.05$, maximum possible vaccination coverage $a = 0.96$, and time that vaccination begins $s = 100$. All rates are in units of years. The function $\rho(t) = a(1 - \exp[\sigma(t - s)])$ is the time-dependent vaccination rate. Total population size in this simulation was 100,000 individuals. Solutions were obtained using the adaptive tau-leaping algorithm [25]. Parameters follow Ferguson et al. [26].

Next, we consider whether these phases can be distinguished in data that can be derived from surveillance. To reflect real-world data collection, we first aggregated simulated cases to annual reporting intervals. Recent theoretical developments predict that the coefficient of variation in the number of cases in a moving window should be a sensitive and specific indicator of the approach to the vaccination threshold [22]. Indeed, calculating the coefficient of variation (CV; the ratio of the standard deviation to the mean) in a moving window over the annualized simulated cases in Figure 1 shows a characteristic dynamical pattern: (i) stationary fluctuations in the *pre-campaign era*, (ii) substantial and sustained rise in CV beginning with the start of the elimination campaign during the *approach to the vaccination threshold*, (iii) erratic fluctuations following the achievement of the vaccination threshold but *preceding elimination*, and (iv) a final decline as the pathogen approaches *complete elimination* (Figure 2).

Of course, if one has perfect case reporting, then it is simple enough to inspect the number of case reports over time to determine that the elimination campaign is continuing to realize results—although it is important to note that this will not tell us if the vaccination threshold has been reached. In order to determine if our method works with incomplete case reporting, we looked at case reports that were binomial samples of 95%, 50%, 10%, and 5% of the true number of cases in each year. Surprisingly, despite the fact that case reports differ by three orders of magnitude and themselves contain no information about the true prevalence in the population, the statistical signals indicating the approach

to the vaccination threshold were practically indistinguishable (Figure 3a). This result shows that it may be possible to monitor the path to elimination by taking new approaches to processing surveillance data, even if they are subject to a significant amount of under-reporting. Details of these simulations are provided in online supplementary material.

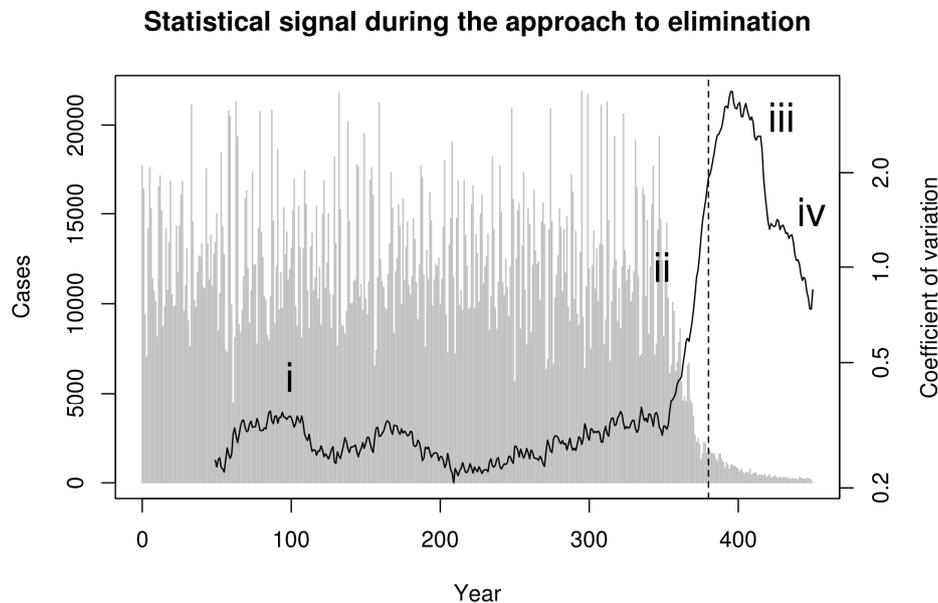


Figure 2. Statistical signal of the approach to the vaccination thresholds in simulated smallpox elimination. Contagion systems exhibit critical slowing down in the approach to a tipping point such as the vaccination threshold. The coefficient of variation in a moving window provides a measurement of the magnitude of this slowing down. We calculated this statistical signature by first detrending with a one-sided filter and then computing the coefficient of variation in a moving window of 30 observations. This figure shows this statistical signal to begin increasing with the onset of vaccination and to rise dramatically as the approach to the vaccination threshold (vertical dashed line) is approached.

These results show that critical slowing down accompanies the approach to the vaccination threshold and is detectable using very simple statistical tests. We also tested if the timing of the vaccination threshold itself is identifiable using these methods. Motivated by the theoretical prediction that the coefficient of variation would diverge at the vaccination threshold [22], we approximated the moving window signal with the hyperbolic equation $w_t = a/(t - c) + b$, where w_t is the calculated value of the signal at time t , a is a fit coefficient that governs the speed of divergence, b is a fit coefficient corresponding to the baseline value of w , and c is the location of the threshold. To test this model, we simulated 1000 realizations of the epidemic model in Figure 1. We then automated the estimation of the threshold by fitting a , b , and c by nonlinear least squares to data for each simulation from its start until the point it reached 75% of its maximum value. The estimated time that the threshold would be crossed was slightly biased and independent of the level of under-reporting (Figure 3b). We suspect the bias evident in Figure 3b is due to the approximation $w_t = a/(t - c) + b$. Accordingly, we consider the observed bias (<1%) to be remarkably small. We consider it a high priority for further work to develop a more complete understanding of the conditions under which the time of threshold crossing can be accurately predicted.

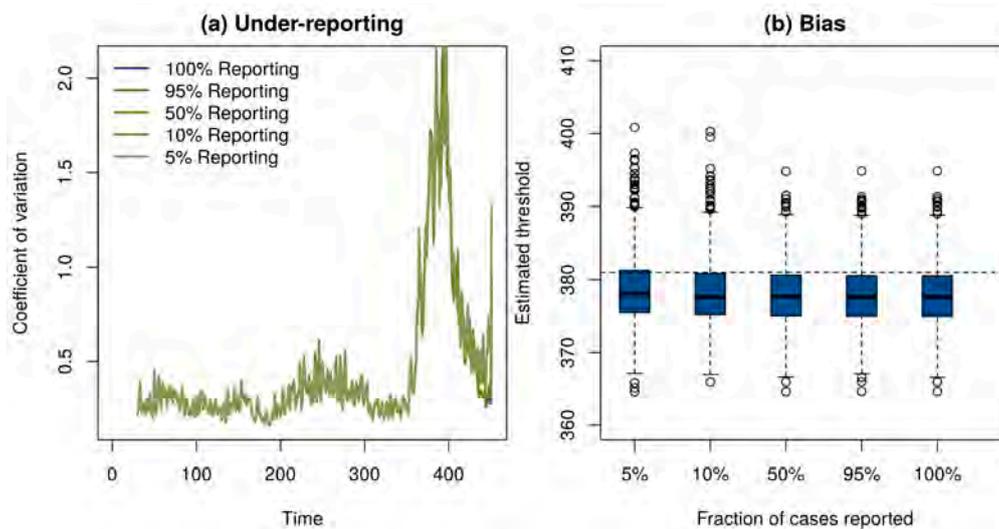


Figure 3. (a) Under-reporting had negligible effect on the ability of the moving window coefficient of variation to respond to the approach to the vaccination threshold. (b) A hyperbolic approximation was used to forecast the time the vaccination threshold would be reached. The estimated crossing time in 1000 simulations (boxplots) was on average slightly earlier than the true time of year 381 (horizontal dashed line). Under-reporting had little effect on the precision or bias of predictions.

3. Conclusions

In conclusion, these examples demonstrate two key properties of the dynamics of an infectious disease on the road to elimination. First, there is an inevitable delay between the time that the elimination threshold is achieved and the time that the last infected individual acquires infection. This is related to the phenomenon of *bifurcation delay*, described by Dibble et al. [27] in the context of disease emergence. Second, the statistical properties of fluctuations in the number of cases can indicate where along the path to elimination a particular population lies. In the approach to the vaccination threshold studied here, the coefficient of variation in a moving window increased continuously with vaccination pressure, providing a possible way to document gains from increases in effort. Areas for further research include determining what statistics—among dozens now available [28]—give both sensitive and specific warning of epidemiological transitions [22,24], and whether or not spatially explicit analogs of the approach taken here may be more powerful [29]. Additionally, it will be important to establish how well these indicators perform under the more complicated conditions that affect real elimination campaigns, including spatial heterogeneity, control programs that may be intermittent or vary in intensity, and seasonal forcing. Moreover, further theoretical research is needed to establish that critical slowing down will be detectable in parasites with more complicated life cycles (i.e., schistosomes and Guinea worm). The evidence so far—primarily from analysis of the near-critical dynamics of the Ross–MacDonald model—suggests that vector-borne diseases like malaria are predicted to display critical slowing down [24].

Regardless of further developments, immediate and simple insights could be derived from simple simulation studies on the frequency and accuracy of reporting needed to monitor the ongoing near-critical contagion systems of polio in places where it remains endemic.

Supplementary Materials: The supplementary materials are available online at <http://www.mdpi.com/2414-6366/2/3/20/s1>.

Acknowledgments: Research reported here was supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award Number U01GM110744. The content is solely the responsibility of the authors and does not necessarily reflect the official views of the National Institutes of Health. S.I.H. is funded by grants from the Bill & Melinda Gates Foundation (OPP1119467, OPP1093011, OPP1106023, and OPP1132415).

Author Contributions: J.M.D. and S.I.H. conceived the idea and designed the study; J.M.D. performed the analysis; J.M.D. and S.I.H. wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest. The founding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, and in the decision to publish the results.

References

1. Hopkins, D.R. Disease eradication. *N. Engl. J. Med.* **2013**, *368*, 54–63. [[CrossRef](#)] [[PubMed](#)]
2. Fenner, F.; Henderson, D.A.; Arita, I.; Jecek, Z.; Ladnyi, I.D. *Smallpox and Its Eradication. History of International Public Health*; World Health Organization: Geneva, Switzerland, 1998.
3. Grassly, N.C. The final stages of the global eradication of poliomyelitis. *Philos. Trans. R. Soc. B* **2013**, *368*, 20120140. [[CrossRef](#)] [[PubMed](#)]
4. Biswas, G.; Sankara, D.P.; Agua-Agum, J.; Maiga, A. Dracunculiasis (guinea worm disease): Eradication without a drug or a vaccine. *Philos. Trans. R. Soc. B* **2013**, *368*, 20120146. [[CrossRef](#)] [[PubMed](#)]
5. Rinaldi, A. Yaws eradication: Facing old problems, raising new hopes. *PLoS Negl. Trop. Dis.* **2012**, *6*, e1837. [[CrossRef](#)] [[PubMed](#)]
6. Global Health Group; Malaria Atlas Project. *Atlas of Malaria Eliminating Countries*; The Global Health Group, Global Health Sciences, University of California: San Francisco, CA, USA, 2011.
7. World Health Organization. Global programme to eliminate lymphatic filariasis: Progress report, 2011. *Wkly. Epidemiol. Rec.* **2012**, *37*, 345–356.
8. Brady, M. Global Alliance to Eliminate Lymphatic Filariasis. Seventh meeting of the Global Alliance to Eliminate Lymphatic Filariasis: Reaching the vision by scaling up, scaling down, and reaching out. *Parasites Vectors* **2014**, *7*, 46. [[CrossRef](#)] [[PubMed](#)]
9. Ferrari, M.J.; Grenfell, B.T.; Strebel, P.M. Think globally, act locally: The role of local demographics and vaccination coverage in the dynamic response of measles infection to control. *Philos. Trans. R. Soc. B* **2013**, *368*, 20120141. [[CrossRef](#)] [[PubMed](#)]
10. Papania, E.J.; Wallace, G.S.; Rota, P.A.; Icenogle, J.P.; Fiebelkorn, A.P.; Armstrong, G.L.; Reef, S.E.; Redd, S.B.; Abernathy, E.S.; Barskey, A.E.; et al. Elimination of endemic measles, rubella, and congenital rubella syndrome from the Western hemisphere: The US experience. *JAMA Pediatr.* **2014**, *168*, 148–155. [[CrossRef](#)] [[PubMed](#)]
11. Mackenzie, C.D.; Homeida, M.M.; Hopkins, A.D.; Lawrence, J.C. Elimination of onchocerciasis from Africa: Possible? *Trends Parasitol.* **2012**, *28*, 16–22. [[CrossRef](#)] [[PubMed](#)]
12. Traore, M.O.; Sarr, M.D.; Badji, A.; Bissan, Y.; Diawara, L.; Doumbia, K.; Goita, S.F.; Konate, L.; Mounkoro, K.; Seck, A.F.; et al. Proof-of-principle of onchocerciasis elimination with ivermectin treatment in endemic foci in Africa: Final results of a study in Mali and Senegal. *PLoS Negl. Trop. Dis.* **2012**, *6*, e1825. [[CrossRef](#)] [[PubMed](#)]
13. Rollinson, D.; Knopp, S.; Levitz, S.; Stothard, J.R.; Tchuente, L.-A.T.; Garba, A.; Mohammed, K.A.; Schur, N.; Person, B.; Colley, D.G.; et al. Time to set the agenda for schistosomiasis elimination. *Acta Trop.* **2013**, *128*, 423–440. [[CrossRef](#)] [[PubMed](#)]
14. Franco, J.R.; Simarro, P.P.; Diarra, A.; Ruiz-Postigo, J.A.; Jannin, J.G. The journey towards elimination of gambiense human African trypanosomiasis: Not far, nor easy. *Parasitology* **2014**, *141*, 748–760. [[CrossRef](#)] [[PubMed](#)]
15. Klepac, P.; Metcalf, J.E.; McLean, A.R.; Hampson, K. Towards the endgame and beyond: Complexities and challenges for the elimination of infectious diseases. *Philos. Trans. R. Soc. B* **2013**, *368*, 20120137. [[CrossRef](#)] [[PubMed](#)]
16. Bousema, T.; Griffin, J.T.; Sauerwein, R.W.; Smith, D.L.; Churcher, T.S.; Takken, W.; Ghani, A.; Drakeley, C.; Gosling, R. Hitting hotspots: Spatial targeting of malaria for control and elimination. *PLoS Med.* **2012**, *9*, e1001165. [[CrossRef](#)] [[PubMed](#)]
17. Cohen, J.M.; Smith, D.L.; Cotter, C.; Ward, A.; Yamey, G.; Sabot, O.J.; Moonen, B. Malaria resurgence: A systematic review and assessment of its causes. *Malar. J.* **2012**, *11*, 122. [[CrossRef](#)] [[PubMed](#)]
18. Cochi, S.L.; Linkins, R.W. The final phase of polio eradication: New vaccines and complex choices. *J. Infect. Dis.* **2012**, *205*, 169–171. [[CrossRef](#)] [[PubMed](#)]

19. Kew, O. Reaching the last one per cent: Progress and challenges in global polio eradication. *Curr. Opin. Virol.* **2012**, *2*, 188–198. [[CrossRef](#)] [[PubMed](#)]
20. Omer, S.B.; Orenstein, W.A.; Koplan, J.P. Go big and go fast—Vaccine refusal and disease eradication. *N. Engl. J. Med.* **2013**, *368*, 1374–1376. [[CrossRef](#)] [[PubMed](#)]
21. Scheffer, M.; Carpenter, S.R.; Lenton, T.M.; Bascompte, J.; Brock, W.; Dakos, V.; van de Koppel, J.; van de Leemput, I.A.; Levin, S.A.; van Nes, E.H.; et al. Anticipating critical transitions. *Science* **2012**, *338*, 344–348. [[CrossRef](#)] [[PubMed](#)]
22. O'Regan, S.M.; Drake, J.M. Theory of early warning signals of disease emergence and leading indicators of elimination. *Theor. Ecol.* **2013**, *6*, 333–357. [[CrossRef](#)]
23. Boettiger, C.; Ross, N.; Hastings, A. Early warning signals: The charted and uncharted territories. *Theor. Ecol.* **2013**, *6*, 255–264. [[CrossRef](#)]
24. O'Regan, S.M.; Lillie, J.W.; Drake, J.M. Leading indicators of mosquito-borne disease elimination. *Theor. Ecol.* **2016**, *9*, 269–286. [[CrossRef](#)] [[PubMed](#)]
25. Cao, Y.; Gillespie, D.T.; Petzold, L.R. Efficient step size selection for the tau-leaping simulation method. *J. Chem. Phys.* **2006**, *124*, 044109. [[CrossRef](#)] [[PubMed](#)]
26. Ferguson, N.M.; Keeling, M.J.; Edmunds, W.J.; Gani, R.; Grenfell, B.T.; Anderson, R.M.; Leach, S. Planning for smallpox outbreaks. *Nature* **2003**, *426*, 681–685. [[CrossRef](#)] [[PubMed](#)]
27. Dibble, C.; O'Dea, E.A.; Park, A.W.; Drake, J.M. Waiting time to infectious disease emergence. *J. R. Soc. Interface* **2016**, *13*, 20160540. [[CrossRef](#)] [[PubMed](#)]
28. Dakos, V.; Carpenter, S.R.; Brock, W.A.; Ellison, A.M.; Guttal, V.; Ives, A.R.; Kéfi, S.; Livina, V.; Seekell, D.; van Nes, E.H.; et al. Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PLoS ONE* **2012**, *7*, e41010. [[CrossRef](#)] [[PubMed](#)]
29. Guttal, V.; Jayaprakash, C. Spatial variance and spatial skewness: Leading indicators of regime shifts in spatial ecological systems. *Theor. Ecol.* **2009**, *2*, 3–12. [[CrossRef](#)]



© 2017 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

Supplementary Materials: Monitoring the Path to the Elimination of Infectious Diseases

John M. Drake and Simon I. Hay

Introduction

This document is the supplementary material to the paper “Monitoring the path to the elimination of infectious diseases”. This document provides R code to reproduce the simulated data and all figures contained in the main text as well as additional information. The thesis of that paper is that the phenomenon of *critical slowing down* may be used to document the elimination of endemic infectious disease through vaccination or other interventions. Here we study the dynamics of a smallpox-like pathogen during a program of vaccine rollout and elimination. We provide methods for estimating early warning signals such as the divergence of variance as the critical threshold is approached and demonstrate their robustness to under-reporting. We then develop a method for estimating the time that the vaccination campaign crosses the critical point and compare the performance of this method with a simple extrapolation strategy. It is shown that critical slowing down provides a robust and accurate way to monitor the elimination of an infectious disease.

Setup

The analysis will require some functionality from Carl Boettiger’s `earlywarning` package, which is available on Github at <https://github.com/cboettig/earlywarning>.

```
> require(earlywarning)
```

To reproduce the published figures, we require the following scheme of named colors.

```
> ose1 <- rgb(85, 108, 17, m = 255)
> ose2 <- rgb(160, 108, 17, m = 255)
> ose3 <- rgb(114, 132, 56, m = 255)
> ose4 <- rgb(137, 152, 87, m = 255)
```

Simulations

We study simulations parameterized to represent smallpox. From Ferguson et al. (Ferguson et al. 2003. *Nature* 426, 681-685) we have that smallpox is SEIR-like with 30% mortality. R_0 is between

4 and 10. There is a 12 day incubation period (E-stage) and 12 day infectious period (I-stage; 3 day prodromal stage and 9 day symptomatic period). If we ignore the E-stage, then recovery rate is about 30.5 per year ($(12/360)^{-1} \approx 30.417$). This yields a vaccination threshold of 0.75 to 0.9, which is also about right.

For the simulated scenario, vaccination rate is defined to be a nonlinear function of time. The following R function returns the vaccination rate for a given time based on provided values for the vaccination start time and speed of rollout.

```
> rho.curve.ramp <- function(t, start = 100, speed = -0.08) {
+   ifelse(t <= start, rho <- 0, rho <- 0.96 * (1 - exp((t -
+     start) * speed)))
+   return(rho)
+ }
```

The direct method (Gillespie) is too slow for these parameters, so we use the `adaptivetau` package for adaptive tau leaping. Possible transitions and propensities (stochastic rates) are encoded as follows.

```
> require(adaptivetau)
> transitions = cbind(c(1, 0, 0), c(-1, 0, 0), c(0, -1, 0), c(0,
+ 0, -1), c(-1, 1, 0), c(0, -1, 1), c(0, 0, 1))
> transitions2 = cbind(c(1, 0, 0, 0), c(-1, 0, 0, 0), c(0, -1,
+ 0, 0), c(0, 0, -1, 0), c(-1, 1, 0, 1), c(0, -1, 1, 0), c(0,
+ 0, 1, 0))
> rates <- function(x, params, t) {
+   return(c(params$alpha + params$mu * (x["X"] + x["Y"] + x["Z"]),
+     params$mu * x["X"], params$mu * x["Y"], params$mu * x["Z"],
+     params$beta * x["X"] * x["Y"] / (x["X"] + x["Y"] + x["Z"]) +
+     params$xi * x["X"], params$gamma * x["Y"]))
+ }
> rates.vacc <- function(x, params, t) {
+   rho <- rho.curve.ramp(t, start = params$rho.start, speed = params$rho.speed)
+   return(c(params$alpha + params$mu * (x["X"] + x["Y"] + x["Z"]) *
+     (1 - rho), params$mu * x["X"], params$mu * x["Y"], params$mu *
+     x["Z"], params$beta * x["X"] * x["Y"] / (x["X"] + x["Y"] +
+     x["Z"]) + params$xi * x["X"], params$gamma * x["Y"],
+     params$mu * (x["X"] + x["Y"] + x["Z"]) * (rho)))
+ }
```

Model parameters are declared, grouped into a vector, and passed to the simulator.

```
> mu <- 1/60
> R0 <- 4
> gamma <- 365/12
> pop.size <- 1e+06
> beta <- R0 * (gamma + mu)
> ee0 <- (mu/beta) * pop.size * (R0 - 1)
> Y0 <- floor(ee0)
> Z0 <- floor(1 * (1 - 1/R0 - (mu/beta) * (R0 - 1)) * pop.size)
```

```

> X0 <- floor(pop.size - Y0 - Z0)
> params = list(gamma = gamma, beta = beta, mu = mu, xi = 0.001,
+   rho.start = 350, rho.speed = -0.05, alpha = 0)
> set.seed(10281979)
> r = ssa.adaptivetau(c(X = X0, Y = Y0, Z = Z0, W = 0), transitions2,
+   rates.vacc, params, tf = 450)

```

The following code produces Figure 1 of the main paper.

```

> par(mar = c(5, 4, 4, 4) + 0.1)
> matplot(r[, "time"], r[, c("Y")], type = "l", xlab = "Time",
+   ylab = "Prevalance (infected individuals)", col = "gray",
+   main = "Simulated smallpox dynamics during elimination")
> abline(h = ee0)
> p.star <- 1 - 1/R0
> t <- seq(0, 450)
> rho <- unlist(lapply(t, rho.curve.ramp, start = params$rho.start,
+   speed = params$rho.speed))
> crit <- t[which.min(rho < p.star)]
> R0v <- R0 * (1 - rho)
> ee <- ((mu/beta) * pop.size * (R0v - 1)) * (((mu/beta) * pop.size *
+   (R0v - 1)) > 0)
> lines(t, ee, col = "gray", lwd = 3)
> abline(v = crit, lty = 2)
> par(new = TRUE)
> plot(t, rho, xlab = "", ylab = "", xlim = par("usr")[1:2], xaxs = "i",
+   ylim = c(0, 1), type = "l", lwd = 3.5, axes = FALSE, col = "dodgerblue4")
> axis(4)
> mtext("Vaccination coverage", side = 4, line = 2.8)

```

Now we will look at some early warning signals. First we extract the number of new cases in each year using the custom function `annualize` to return periodically aggregated cases. The function `pad.signal` adds NA values as needed so we can use plotting functionality from the `earlywarning` package. The function `plot.signal` will produce the plot we want. Then we calculate moving window statistics across the “observed” epidemic time series. The timing of the critical point is shown with a vertical line.

```

> annualize <- function(r, period = 1) {
+   r <- data.frame(r)
+   r$year <- r$time%%1
+   r$V <- c(0, diff(r$W))
+   data.annualized <- aggregate(r$V, by = list(r$year), FUN = sum)
+   names(data.annualized) <- c("year", "cases")
+   return(data.annualized)
+ }
> data.annual <- annualize(r)
> pad.signal <- function(x, ws) c(rep(NA, (ws - 1)), x)
> plot.signal <- function(data, crit = NA, ws1 = 6, ws2 = 30, main = "",
+   axes = TRUE, xlab = "Time", ylab = "Coefficient of variation",

```

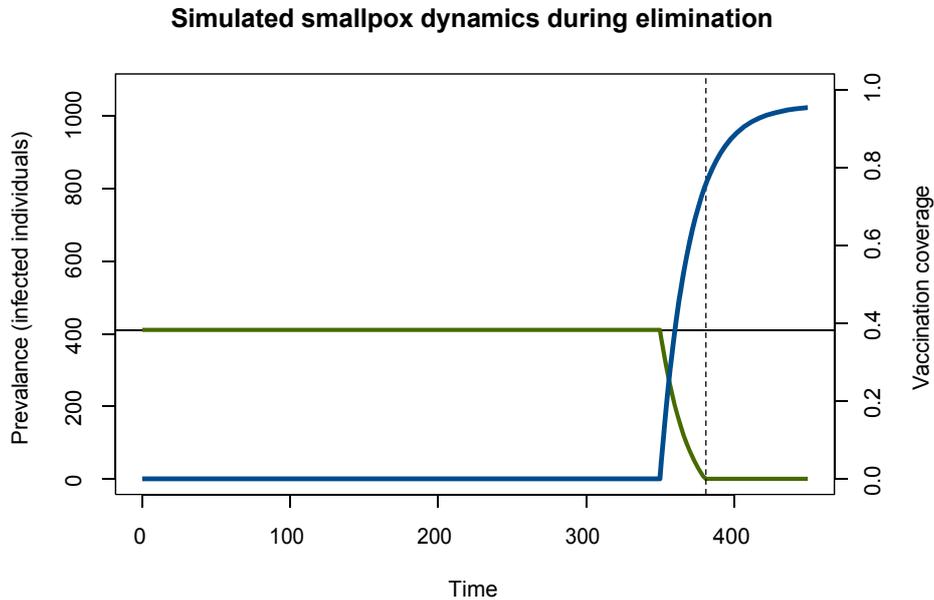


Figure 1: Simulation of a smallpox elimination campaign.

```

+   plot = TRUE) {
+   f <- exp(seq(1, ws1) * (log(0.01/1)/ws1))/sum(exp(seq(1,
+     ws1) * (log(0.01/1)/ws1)))
+   W1 <- filter(data[, 2], filter = f, sides = 1)
+   W2 <- pad.signal(sqrt(window_var(data[, 2], windowsize = ws2)),
+     ws2)/W1
+   if (plot == TRUE) {
+     plot(W2, log = "y", type = "l", xlab = xlab, ylab = ylab,
+       axes = axes, main = main)
+     abline(v = crit, lty = 2)
+   }
+   return(W2)
+ }
> par(mar = c(5, 4, 4, 4) + 0.1)
> plot(data.annual$year, data.annual$cases, type = "h", col = "grey",
+   ylab = "Cases", xlab = "Year")
> par(new = TRUE)
> ews <- plot.signal(data.annual, crit = crit, ws1 = 30, ws2 = 50,
+   main = "Statistical signal during the approach to elimination",
+   xlab = "", ylab = "", axes = FALSE)
> axis(4)
> mtext("Coefficient of variation", side = 4, line = 2.8)
> text(100, 0.5, labels = "i", cex = 1.8)
> text(350, 1, labels = "ii", cex = 1.8)
> text(430, 2.3, labels = "iii", cex = 1.8)

```

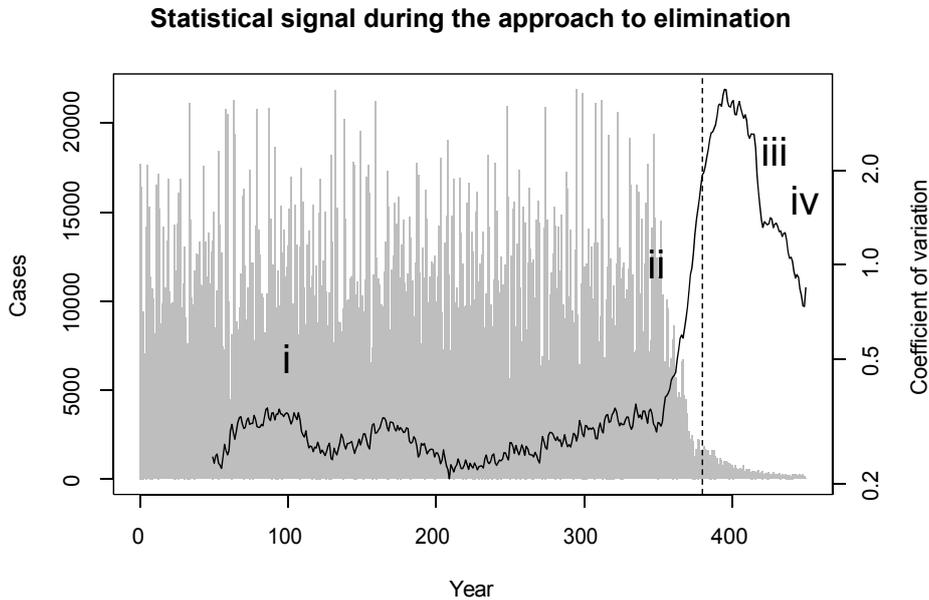


Figure 2: Early warning signals.

```
> text(450, 1.6, labels = "iv", cex = 1.8)
```

Here we look at how the signal is corrupted by binomially sampling the annualized incidence to represent underreporting. Evidently, the signal is quite robust to under-reporting.

```
> par(mfrow = c(5, 1))
> ews <- plot.signal(data.annual, crit = crit, ws1 = 30, ws2 = 50,
+   main = "Statistical signal during the approach to elimination")
> obs <- function(data, p = 0.95) data.frame(time = data$year,
+   cases = rbinom(seq(1, dim(data)[1]), size = data$cases, prob = p))
> ews95 <- plot.signal(w95 <- obs(data.annual, p = 0.95), crit = crit,
+   ws1 = 30, ws2 = 50, main = "95% Case Reporting")
> ews50 <- plot.signal(w50 <- obs(data.annual, p = 0.5), crit = crit,
+   ws1 = 30, ws2 = 50, main = "50% Case Reporting")
> ews10 <- plot.signal(w10 <- obs(data.annual, p = 0.1), crit = crit,
+   ws1 = 30, ws2 = 50, main = "10% Case Reporting")
> ews05 <- plot.signal(w01 <- obs(data.annual, p = 0.05), crit = crit,
+   ws1 = 30, ws2 = 50, main = "5% Case Reporting")
```

Estimating the tipping point

Finally, we look at some ways to estimate the threshold time. First, we seek a transformation to fit a hyperbolic model with linear regression. For instance, assuming the model $f_0 = y = -a/(a - c)$ and taking the inverse we have a linear equation $z = 1/y = (c - x)/a = c/a - x/a = -(1/a)x + c/a$.

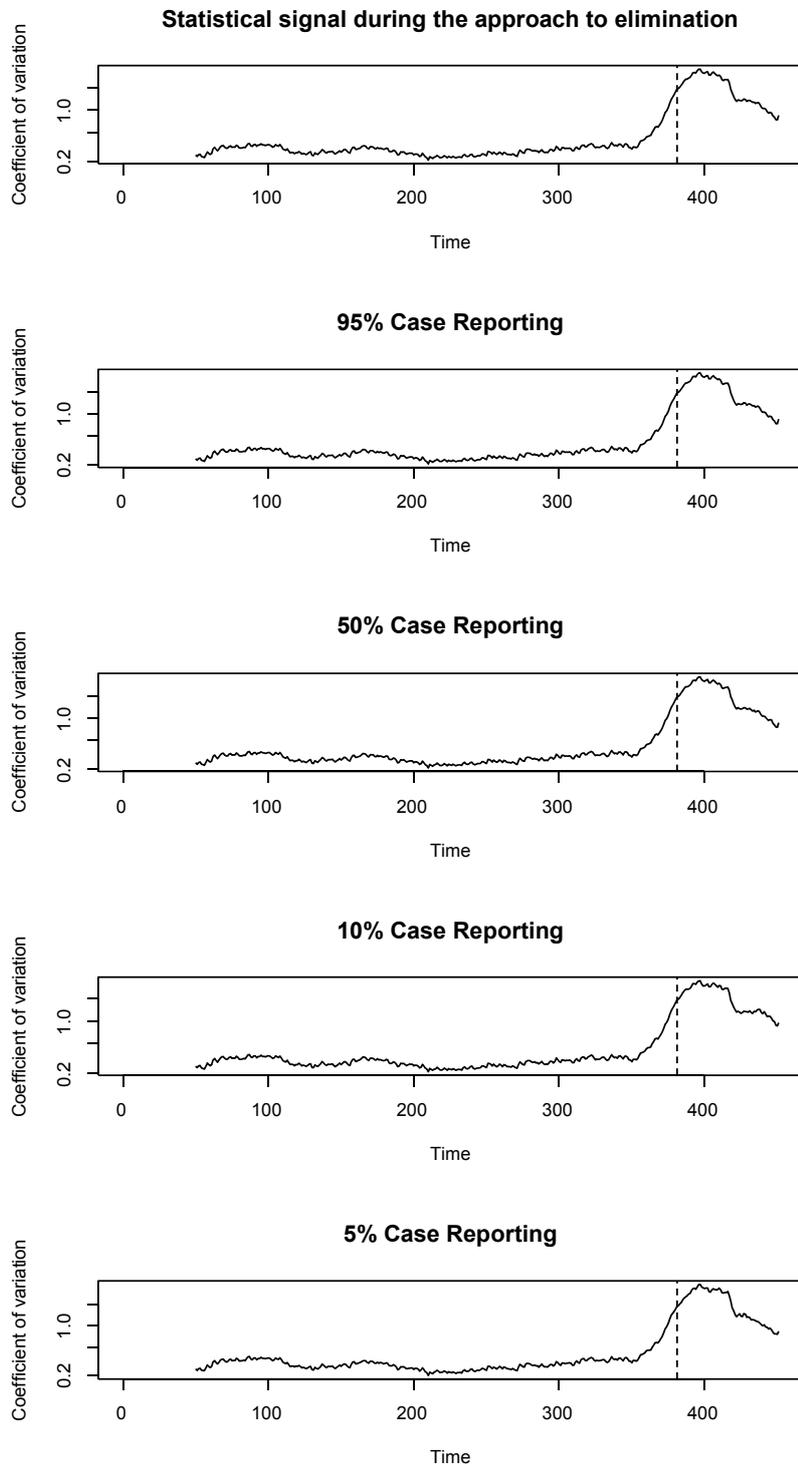


Figure 3: Early warning signals in time series corrupted by sampling.

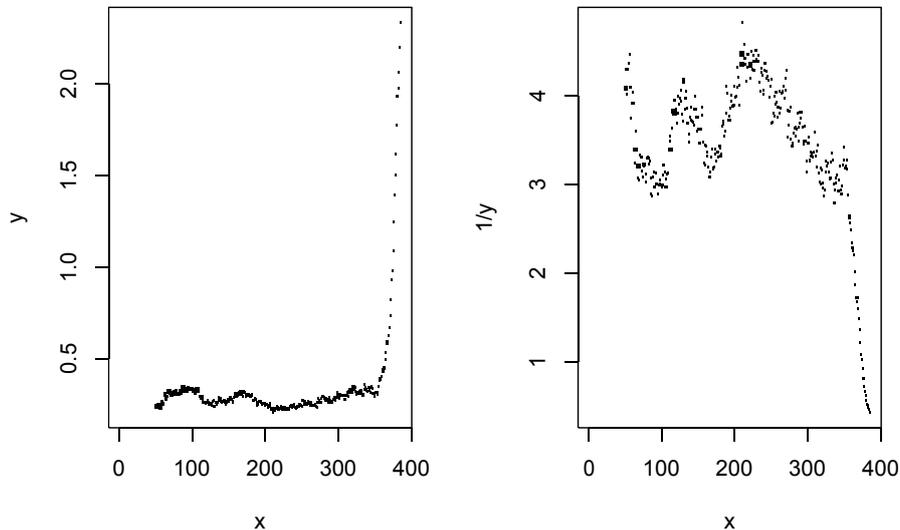


Figure 4: Transformation of a hyperbola model for the early warning signal.

Thus, taking the reciprocal of y , and regressing on x , the slope is the negative inverse of a and the intercept is c/a so that an estimator for the critical time is $c = -\text{slope} \times \text{intercept}$. We try that with the early warning signal computed above.

```
> y <- ews[1:385]
> x <- 1:length(y)
> par(mfrow = c(1, 2))
> plot(x, y, type = "p", pch = ".")
> plot(x, 1/y, type = "p", pch = ".")
> z <- 1/y
> mod <- lm(z[200:385] ~ x[200:385])
> est <- -prod(coef(mod))
> print(est)
```

```
[1]0.132874
```

From Figure 4 we see that at least one problem with this approach is that the rectangular hyperbola $y = a/(x - c)$ assumes an asymptote at zero, which is a reasonable approximation on the original scale but not the inverse scale. Adding an asymptote we have $f_1 = y = a/(x - c) + b$, but this doesn't admit any tidy transformations.

Therefore, we switch to estimating the nonlinear function by least squares

```
> y <- ews[1:365]
> x <- 0:(length(y) - 1)
> plot(x, y)
> mod2 <- nls(y ~ -a/(x - c) + b, start = list(a = 10, b = 0.2,
```

```
+ c = 400))
> print(mod2)
```

```
Nonlinear regression model
model: y ~ -a/(x - c) + b
data: parent.frame()
      a      b      c
1.3371 0.2677 368.8403
residual sum-of-squares: 0.3274
```

```
Number of iterations to convergence: 7
Achieved convergence tolerance: 7.893e-06
```

```
> mod2.refit <- nls(y ~ -a/(x - c) + b, start = coef(mod2))
> print(mod2.refit)
```

```
Nonlinear regression model
model: y ~ -a/(x - c) + b
data: parent.frame()
      a      b      c
1.3371 0.2677 368.8403
residual sum-of-squares: 0.3274
```

```
Number of iterations to convergence: 0
Achieved convergence tolerance: 7.893e-06
```

```
> lines(x, -coef(mod2)[1]/(x - coef(mod2)[3]) + coef(mod2)[2])
```

Figure 5 shows this approach to work much better. In this example, the estimated critical time is 370 compared with a true value of 381. A weakness of this approach is that it can be expected to be numerically unstable so that if insufficient data are available or if start values for the fitting algorithm are poorly chosen it will run into difficulties.

The following function attempts to estimate the critical time directly from a signal such as that generated above.

```
> estimate.threshold <- function(ews, sig.time = NA, plot = FALSE,
+   a = 7.5) {
+   if(is.na(sig.time)) {
+     stop.50 <- 0.75 * max(ews, na.rm = TRUE)
+     sig.time <- min(which(ews > stop.50))
+   }
+   y <- ews[1:sig.time]
+   x <- 0:(length(y) - 1)
+   mod <- nls(y ~ -a/(x - c) + b, start = list(a = a, b = min(ews,
+     na.rm = TRUE), c = 1.05 * sig.time), control = c(max.iter = 500,
+     minFactor = 1/4096))
+   if(plot == TRUE) {
+     plot(x, y, pch = 16, cex = 0.5, col = "grey", xlab = "Time",
+       ylab = "Early warning signal", main = paste("Estimated threshold:"),
```

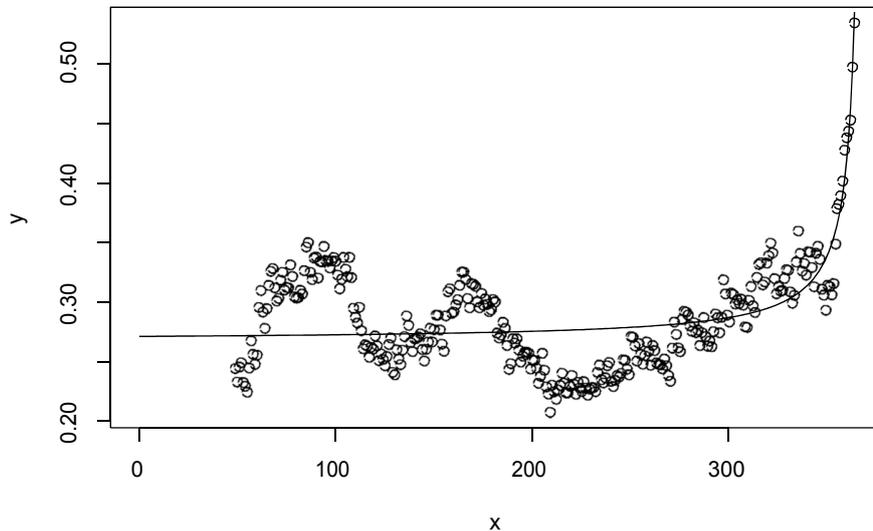


Figure 5: Nonlinear least squares fit to a hyperbola model for the early warning signal.

```
+           round(unnamed(coef(mod))[3]))
+     lines(x, -coef(mod)[1]/(x - coef(mod)[3]) + coef(mod)[2],
+           lwd = 2, col = "dodgerblue4")
+   }
+   return(estimated.threshold <- unnamed(coef(mod))[3])
+ }
```

Its use is demonstrated as follows:

```
> threshold <- estimate.threshold(ews, plot = TRUE)
```

The result is shown in Figure 6.

What might we compare these estimates to? One proposal is simple extrapolation from the observed trend. Of course, this doesn't predict the tipping point, just the extinction time. But, if bifurcation delay is small then these should be close. Our extrapolating estimator will perform the following steps:

1. Select that portion of the incidence data from the last the number of cases was greater than 90% of its maximum to some later time at which the extrapolation is to be made.
2. Log transform the data and fit a linear equation: $\log(w + 1) = mx + b$, where w is annual incidence (one is added to observations particularly to facilitate studies of under-reporting, which will often yield values of zero).
3. Solve the resulting equation for $\log(w + 1) = 0$ yielding estimator $\hat{t} = -b/m$ for the critical time.

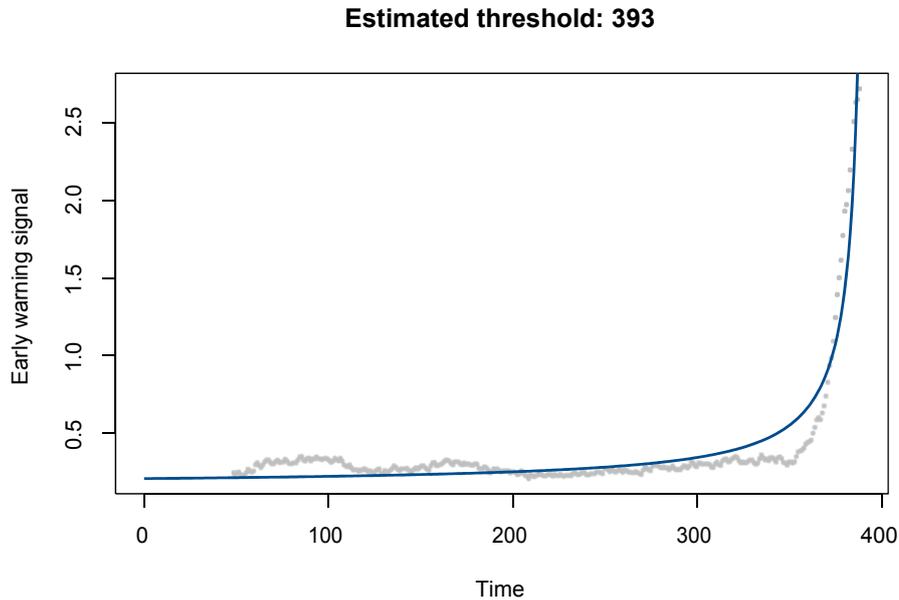


Figure 6: Demonstration of the use of the threshold estimator.

```

> extrapolator <- function(data, sig.time = NA, plot = FALSE) {
+   if (is.na(sig.time))
+     sig.time <- length(data)
+   ma <- max(data, na.rm = TRUE)
+   ma90 <- 0.9 * ma
+   w <- max(which(data > ma90))
+   y <- data[w:sig.time]
+   x <- (w - 1):(sig.time - 1)
+   z <- log(y + 1)
+   mod <- lm(z ~ x)
+   if (plot) {
+     plot(x, z, xlab = "Time", ylab = "Cases", main = paste("Estimated threshold:",
+       round(unname(-coef(mod)[1]/coef(mod)[2])))
+     abline(mod, lwd = 2, col = "dodgerblue4")
+   }
+   return(extrapolated.threshold <- unname(-coef(mod)[1]/coef(mod)[2]))
+ }

```

This function is demonstrated as follows and illustrated in Figure 7

```

> extrapolated.threshold <- extrapolator(data.annual$cases, plot = TRUE)

```

Now, we'd like to see how the bias and variance of the the early warning statistic and extrapolator estimate change with under-reporting in $n = 1000$ simulations.

```

> n.sims <- 1000

```

Estimated threshold: 562

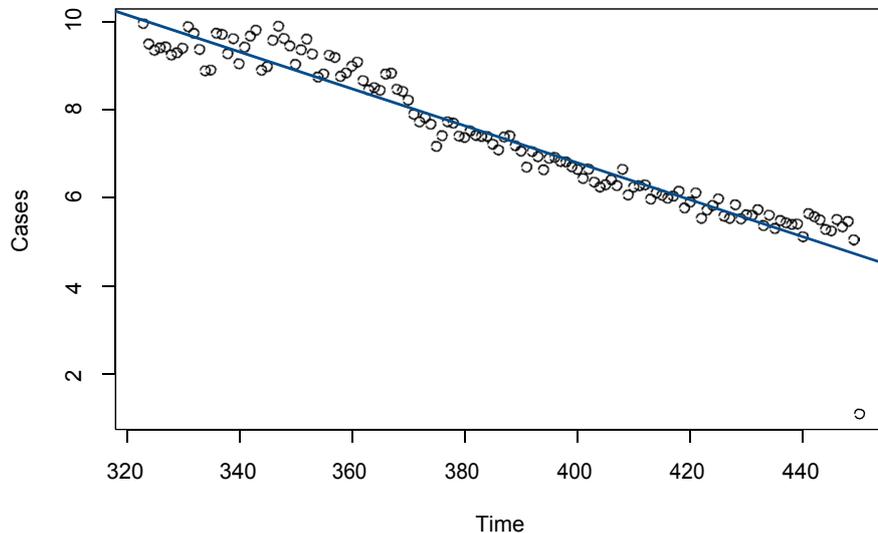


Figure 7: Demonstration of the use of an extrapolation approach to estimating the critical threshold.

```
> sims <- lapply(1:n.sims, function(x) ssa.adaptivetau(c(X = X0,
+   Y = Y0, Z = Z0, W = 0), transitions2, rates.vacc, params,
+   tf = 450))
> save(sims, n.sims, file = "sims.RData")

> load("sims.RData")
> sims.annualized <- rapply(sims, f = function(x) annualize(x),
+   how = "list")
> set.seed(10281979)
> sims95 <- lapply(sims.annualized, function(x) obs(x, p = 0.95))
> sims50 <- lapply(sims.annualized, function(x) obs(x, p = 0.5))
> sims10 <- lapply(sims.annualized, function(x) obs(x, p = 0.1))
> sims05 <- lapply(sims.annualized, function(x) obs(x, p = 0.05))
> thresh <- unlist(lapply(sims.annualized, FUN = function(x) extrapolator(x$cases)))
> thresh95 <- unlist(lapply(sims95, FUN = function(x) extrapolator(x$cases)))
> thresh50 <- unlist(lapply(sims50, FUN = function(x) extrapolator(x$cases)))
> thresh10 <- unlist(lapply(sims10, FUN = function(x) extrapolator(x$cases)))
> thresh05 <- unlist(lapply(sims05, FUN = function(x) extrapolator(x$cases)))
> sims.ews <- lapply(sims.annualized, FUN = function(x) plot.signal(x,
+   plot = FALSE))
> sims95.ews <- lapply(sims95, FUN = function(x) plot.signal(x,
+   plot = FALSE))
> sims50.ews <- lapply(sims50, FUN = function(x) plot.signal(x,
+   plot = FALSE))
> sims10.ews <- lapply(sims10, FUN = function(x) plot.signal(x,
```

```

+   plot = FALSE))
> sims05.ews <- lapply(sims05, FUN = function(x) plot.signal(x,
+   plot = FALSE))
> thresh.ews <- unlist(lapply(sims.ews, FUN = function(x) estimate.threshold(x)))
> thresh95.ews <- unlist(lapply(sims95.ews, FUN = function(x) estimate.threshold(x)))
> thresh50.ews <- unlist(lapply(sims50.ews, FUN = function(x) estimate.threshold(x)))
> thresh10.ews <- unlist(lapply(sims10.ews, FUN = function(x) estimate.threshold(x)))
> thresh05.ews <- unlist(lapply(sims05.ews, FUN = function(x) estimate.threshold(x)))
> threshold <- data.frame(threshold = c(thresh.ews, thresh95.ews,
+   thresh50.ews, thresh10.ews, thresh05.ews, thresh, thresh95,
+   thresh50, thresh10, thresh05), reporting = rep(c(rep(100,
+   n.sims), rep(95, n.sims), rep(50, n.sims), rep(10, n.sims),
+   rep(5, n.sims)), 2), method = c(rep("ews", n.sims * 5), rep("extrapolation",
+   n.sims * 5)))
> boxplot(threshold ~ method + reporting, data = threshold, col = rep(c("dodgerblue4",
+   ose1), 5), axes = FALSE, log = "", ylim = c(350, 1500))
> axis(2)
> axis(1, at = seq(0, 8, by = 2) + 1.5, labels = c("5%", "10%",
+   "50%", "95%", "100%"))
> box()
> legend("topleft", bty = "n", col = c("dodgerblue4", ose1), legend = c("Critical slowing down
+   "Extrapolation"), lty = 1, lwd = 6)
> save(sims.annualized, sims95, sims50, sims10, sims05, thresh,
+   thresh95, thresh50, thresh10, thresh05, sims.ews, sims95.ews,
+   sims50.ews, sims10.ews, sims05.ews, thresh.ews, thresh95.ews,
+   thresh50.ews, thresh10.ews, thresh05.ews, threshold, file = "bias.RData")

```

Of course, since the extrapolator has no information about bifurcation delay it will inevitably be very biased. What's more, there's no "natural" way to correct for this bias. Therefore, we will draw a plot showing just the effect of under-reporting on the critical slowing down based estimator.

```

> load("bias.RData")
> par(mfrow = c(1, 2))
> par(mar = c(5, 4, 2, 0) + 0.1)
> plot(sims.ews[[1]], lwd = 2, col = "dodgerblue4", ylab = "Coefficient of variation",
+   main = "(a) Under-reporting")
> lines(sims95.ews[[1]], lwd = 2, col = ose1)
> lines(sims50.ews[[1]], lwd = 2, col = ose2)
> lines(sims10.ews[[1]], lwd = 2, col = ose3)
> lines(sims05.ews[[1]], lwd = 2, col = ose4)
> legend("topleft", legend = c("100% Reporting", "95% Reporting",
+   "50% Reporting", "10% Reporting", "5% Reporting"), col = c("dodgerblue4",
+   ose1, ose2, ose3, ose4), lwd = 2, bty = "n")
> boxplot(threshold ~ method + reporting, data = threshold[threshold$method ==
+   "ews", ], col = rep(c("dodgerblue4", ose1), 5), axes = FALSE,
+   ylim = c(360, 410), xlab = "Fraction of cases reported",
+   ylab = "Estimated threshold", main = "(b) Bias")
> axis(2)

```

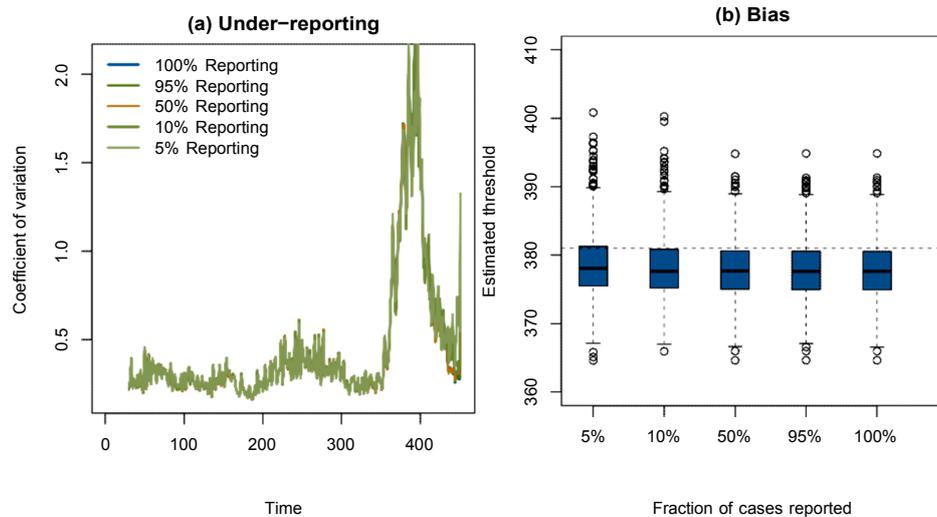



Figure 9: Bias in an approach to estimating the critical threshold from measurement of critical slowing down. Note that lines in the left panel are largely, but not exactly, overlapping.

```
> axis(1, at = seq(0, 8, by = 2) + 1, labels = c("5%", "10%", "50%",
+ "95%", "100%"))
> box()
> abline(h = crit, lty = 2)
```

Figure 9 shows that there is a small amount of negative bias in this critical slowing down estimator of the critical point (the threshold is predicted to be earlier than than it actually is). One thing I've found is that this is somewhat sensitive to the time at which the prediction is made. Possibly this is also due to the size of the moving window used for calculating the variance. In this experiment this is set to 30 observations. A smaller moving window would have less inertia (would adjust faster), but preliminary numerical experiments found this to trade off with the estimability of a in the hyperbola model: with fewer than about 30 observations in the window it was not possible to find starting positions for a that would actually converge. Some additional experiments showed that models could be fit with reporting as low as 1% or 2% but automated fitting was impossible and the variance in the estimate became very large (though not as large as for the extrapolation method).