This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

Robust Economic Model Predictive Control of Continuous-time Epidemic Processes

Nicholas J. Watkins, Cameron Nowzari, and George J. Pappas

***Abstract*—In this paper, we develop a robust economic model predictive controller for the containment of stochastic Susceptible-Exposed-Infected-Vigilant $(SEIV)$ epidemic processes which drives the process to extinction quickly, while minimizing the rate at which control resources are used. The work we present here is significant in that it addresses the problem of efficiently controlling general stochastic epidemic systems without relying on mean-field approximation, which is an important issue in the theory of stochastic epidemic processes. This enables us to provide rigorous convergence guarantees on the stochastic epidemic model itself, improving over the mean-field type convergence results of most prior work. There are two primary technical difficulties addressed in treating this problem: (i) constructing a means of tractably approximating the evolution of the process, so that the designed approximation is robust to the modeling error introduced by the applied moment closure, and (ii) guaranteeing that the designed controller causes the closed-loop system to drive the $SEIV$ process to extinction quickly. As an application, we use the developed framework for optimizing the use of quarantines in containing an $SEIV$ epidemic outbreak.**

***Index Terms*—Epidemic Processes, Model Predictive Control, Networked Systems.**

I. Introduction

The modern study of epidemic models has been intense for the past several decades, with work dating back to the 1970s [1], and a slew of recent results from the control community coming in recent years [2]–[14]. Potential applications include message passing in complex wireless networks [5], competition between multiple mimetic behaviors in social networks [15], and the spread of biological disease [16]–[19]. A recent review of the control of epidemic processes can be found in [20].

An important problem in the field is understanding how to efficiently control epidemic processes in such a manner so as to be able to provide rigorous performance guarantees on the statistics of the process. This problem arises from the complexity inherent in networked epidemic systems: the stochastic dynamics which describe the fundamental aspects of the process entangle the components of the system's state, making their analysis inherently difficult. For sufficiently simple epidemics, such as the Susceptible-Infected-Susceptible (SIS) process, using a mean-field type moment closure, in which the second-order moments of the system are approximated by products of first-order moments, yields dynamics which provide an upper-bound for the expectation of the stochastic process [21]. In general, this is not the case. Indeed, even for simple models with multiple compartments, simulations have shown standard mean-field approximations to be unreliable proxies for the statistics of the underlying stochastic process (see, e.g., [22]). As such, it remains a scientifically interesting question to develop control techniques for stochastic epidemic networked processes for which we can make rigorous claims about the statistics of the process. As the types of controls available to authorities for the prevention of disease are often costly, understanding well how to optimize control resource use while still guaranteeing that the epidemic will end quickly is of critical importance in mitigating the effect of future epidemic threats, such as the predicted increase in endemic diseases due to climate change [23], or the emergence of drug-resistant superbugs [24], [25].

There are two approaches considered in prior work concerning the control of stochastic epidemic models without mean-field approximation. In one approach [12], [13], authors study policies which vary the healing rates of the nodes in the graph according to some computed priority order, wherein infected nodes with higher priority are treated before others. These works guarantee that if sufficient healing resources are available, such a policy can control the SIS epidemic to the disease-free state quickly. In the other approach [14], a model predictive controller is developed which controls the statistics of a discrete-time SIS epidemic directly by varying the processes' spreading parameters. As model predictive control is an important paradigm in control theory with many successful applications in diverse fields (see, e.g. [26]–[29] and the references therein), and it is unclear if the priority-order strategies of [12], [13] can be extended to more general settings efficiently, we work on generalizing the approach of [14] here to the setting of continuous-time epidemics.

The primary contribution of this paper is the development of a robust economic model predictive control scheme for the stochastic continuous-time $SEIV$ process which allocates control resources to the spreading network so as to minimize the resource consumption rate realized by the controller, while providing a strong theoretical convergence guarantee that the disease will be eliminated from the network quickly. This is the first paper to use model predictive control in the context of epidemic containment for continuous-time epidemic processes, and one of the first papers to develop feedback controls for continuous-time stochastic epidemic processes. With respect to our earlier work [14], the text

N.J. Watkins and G.J. Pappas are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA, {nwatk,pappasg}@upenn.edu. C. Nowzari is with the Department of Electrical and Computer Engineering, George Mason University, Fairfax, VA 22030, USA, cnowzari@gmu.edu.

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

presented here differs in that it considers the control of continuous-time epidemic processes, wherein propagating the uncertainty of the system is inherently difficult, and rigorous arguments concerning the convergence of the closed-loop system are more difficult to develop. The two key technical difficulties encountered in addressing this task are (i) providing a rigorous moment-closure approximation of the $SEIV$ process which provides bounds of its statistics, and (ii) providing a convergence analysis of the process when evolving under the designed model predictive controller which guarantees that the amount of time that passes until the process is disease-free is small.

It is important to underline that we believe the closure techniques presented here will generalize to further epidemic models readily, as will the convergence arguments used. We expect that the framework presented here will provide a fertile avenue for future research, providing a framework in which problems studied in prior works considering only the mean-field regime (such as [2]–[11]) can be readily extended.

Organization of Remainder: The remainder of the paper is organized as follows. In Section II, we detail the $SEIV$ epidemic model, the control architecture, and the problem statement. In Section III, we develop a moment closure for the $SEIV$ epidemic model which provides the robust approximation guarantees which are required for the implementation of the desired control scheme. In Section IV, we perform a convergence analysis of the developed controller. In Section V, we apply the developed controller to a problem of minimizing the number of nodes put into quarantine in order to guarantee a specified exponential decay of the number of exposed and infected nodes in the system. In Section VI, we summarize the main results of the paper, and comment on avenues for future work. •

Notation and Terminology: We use $\mathbb{R}_{\geqslant 0}$ to denote the set of non-negative real numbers, and $\mathbb{Z}_{\geqslant 0}$ denote the set of non-negative integers. We denote by $[k]$ the set of the first k positive integers, i.e. $[k] \triangleq \{1, 2, \dots, k\}$, and by $[k]_0$ the first $k+1$ natural numbers, i.e. $[k]_0 \triangleq \{0, 1, 2, \dots, k\}$. We let e_i denote the i'th column of an identity matrix, with the appropriate dimension inferred from context.

We denote by $\mathbb{E}[X]$ the expectation of a random variable X. Note that when the measure of the expectation is clear from context, we omit it. When necessary, we explicitly include it as a subscript of the operator, i.e. $\mathbb{E}_{\mu}[X]$ is the expectation of X with respect to the measure μ. When clear from context, we omit the initial condition $X(0)$ of a stochastic process. •

II. Model and Problem Statement

In this section, we formally develop the model and the problem we study in this paper. The particular construction of the $SEIV$ process we present here is our own, however the final system we arrive at is the same $SEIV$ model as studied in prior work (see, e.g., [3], [30]). Note that while epidemic models are inherently general mathematical abstractions which may be studied in a variety of contexts, the language we use throughout the paper is made specific to the context of biological epidemics for simplicity.

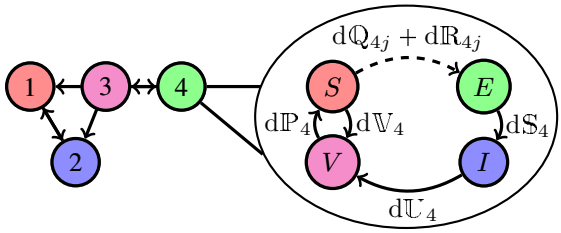


Fig. 1: A compartmental diagram of the $SEIV$ process. The transition process for node 4 is explicitly illustrated, where the measures of the contact processes are included to indicate which process determines which transition, as described in Section II-A.

A. Susceptible-Exposed-Infected-Vigilant Model

We consider the dynamics of a general Susceptible-Exposed-Infected-Vigilant ($SEIV$) epidemic model. In this model, each agent in a population is represented in a directed n-node spreading graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ by a particular node $i \in \mathcal{V}$. At each time in the process, every node belongs to one of a set of the model's *compartments*. On an intuitive level, each compartment represents the stage of infection in which the agent currently resides. In the $SEIV$ model, there are four types of model compartments: susceptible (denoted by the symbol S), exposed (denoted by the symbol E), infected (denoted by the symbol I), and vigilant (denoted by the symbol V). When an agent is susceptible, we may think of it as healthy. When an agent is exposed, we may think of it as having recently come into contact with a contagious disease, but without having yet outwardly displayed symptoms. When an agent is infected, we may think of it as symptomatic. When an agent is vigilant, we may think of it as actively protecting itself against exposure to the disease.

We denote by $X(t)$ a stochastic vector containing the compartmental memberships of each node at time t. To make the notation as intuitive as possible, we index $X(t)$ in two dimensions: one which indicates the compartment which is being described, and the other the numerical label of the node. As such, we denote by $X_i^C(t)$ an indicator random variable, taking the value 1 if node i is in compartment C, where C is one of the symbols $C \in \mathcal{L} \triangleq \{S, E, I, V\}$, and 0 otherwise. In this way, we see that for all times t, we have that $\sum_{C \in \mathcal{L}} X_i^C(t) = 1$ for all i, as each node belongs to precisely one compartment at all times. We denote by $\mathcal{X}$ the set of all possible states of the $SEIV$ process.

Figure 1 provides an illustration of the $SEIV$ process, which can be posed as a system of Itô integrals taken with

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

respect to measures of independent Poisson processes as

$$\begin{aligned}
\mathrm{d}X_i^S &= X_i^V \mathrm{d}\mathbb{P}_i - X_i^S \mathrm{d}\mathbb{V}_i \\
&\quad - \sum_{j\in\mathcal{N}_i} (X_i^S X_j^E \mathrm{d}\mathbb{Q}_{ij} + X_i^S X_j^I \mathrm{d}\mathbb{R}_{ij}), \\
\mathrm{d}X_i^E &= \sum_{j\in\mathcal{N}_i} (X_i^S X_j^E \mathrm{d}\mathbb{Q}_{ij} + X_i^S X_j^I \mathrm{d}\mathbb{R}_{ij}) - X_i^E \mathrm{d}\mathbb{S}_i, \quad (1) \\
\mathrm{d}X_i^I &= X_i^E \mathrm{d}\mathbb{S}_i - X_i^I \mathrm{d}\mathbb{U}_i, \\
\mathrm{d}X_i^V &= X_i^I \mathrm{d}\mathbb{U}_i + X_i^S \mathrm{d}\mathbb{V}_i - X_i^V \mathrm{d}\mathbb{P}_i,
\end{aligned}$$

where $\mathcal{N}_i$ is taken to be the set of in-neighbors of node i on $\mathcal{G}$, $\mathrm{d}\mathbb{P}_i$ is the probability measure induced by the Poisson process generating transition events which take node i from vigilant to susceptible, $\mathrm{d}\mathbb{Q}_{ij}$ is the probability measure induced by the Poisson process generating transition events which take node i from susceptible to exposed through contact with node j when j is exposed, $\mathrm{d}\mathbb{R}_{ij}$ is the probability measure induced by the Poisson process generating transition events which take node i from susceptible to exposed through contact with node j when j is infected, $\mathrm{d}\mathbb{S}_i$ is the probability measure induced by the Poisson process which generates transition events which take node i from exposed to infected, $\mathrm{d}\mathbb{U}_i$ is the probability measure induced by the Poisson process generating transition events which transition node i from infected to vigilant, and $\mathrm{d}\mathbb{V}_i$ is the probability measure associated to the Poisson process which generates transitions from susceptible to vigilant.

We wish to control the process detailed by (1) through the dynamics of its expectation. As such, we must formally derive the expectation dynamics of (1), which after some technical arguments (see Appendix A) can be shown to be

$$\begin{aligned}
\frac{\mathrm{d}\mathbb{E}[X_i^S]}{\mathrm{dt}} &= \alpha_i \mathbb{E}[X_i^V] - \xi_i \mathbb{E}[X_i^S] \\
&\quad - \sum_{j\in\mathcal{N}_i} (\beta_{ij} \mathbb{E}[X_i^S X_j^E] + \gamma_{ij} \mathbb{E}[X_i^S X_j^I]), \\
\frac{\mathrm{d}\mathbb{E}[X_i^E]}{\mathrm{dt}} &= \sum_{j\in\mathcal{N}_i} (\beta_{ij} \mathbb{E}[X_i^S X_j^E] + \gamma_{ij} \mathbb{E}[X_i^S X_j^I]) \\
&\quad - \delta_i \mathbb{E}[X_i^E], \\
\frac{\mathrm{d}\mathbb{E}[X_i^I]}{\mathrm{dt}} &= \delta_i \mathbb{E}[X_i^E] - \eta_i \mathbb{E}[X_i^I], \\
\frac{\mathrm{d}\mathbb{E}[X_i^V]}{\mathrm{dt}} &= \eta_i \mathbb{E}[X_i^I] + \xi_i \mathbb{E}[X_i^S] - \alpha_i \mathbb{E}[X_i^V],
\end{aligned} \quad (2)$$

where α_i, β_{ij}, γ_{ij}, δ_i, η_i, and ξ_i, are the rates of the processes associated to $\mathrm{d}\mathbb{P}_i$, $\mathrm{d}\mathbb{Q}_{ij}$, $\mathrm{d}\mathbb{R}_{ij}$, $\mathrm{d}\mathbb{S}_i$, $\mathrm{d}\mathbb{U}_i$, and $\mathrm{d}\mathbb{V}_i$ respectively, and are termed the *spreading parameters*.

The dynamics (2) are a central object of study in this paper. In particular, a major difficulty in typical epidemic control problems is in reconciling the fact that the induced first-moment dynamics (2) are not closed, as they rely explicitly on the second-order moments $\mathbb{E}[X_i^S X_j^E]$ and $\mathbb{E}[X_i^S X_j^I]$, for which we do not have explicit dynamics. Moreover, it can be shown that the dynamics of the second-order moments depend on third-order moments, and so on until the dynamics entail the expectation of all 4^n possible combinations of node-compartment states (see, e.g. the construction in [31, Section 5]). Because of this, there is no known technique for propagating the exact expectations of the $SEIV$ process forward in time tractably, and techniques for closing (2) by approximating the second-order moment terms - called *moment closures* - are typically employed. One contribution of this work is a novel moment closure which provides over- and under- approximations of compartmental membership probabilities for each node and every compartment, presented in Section III.

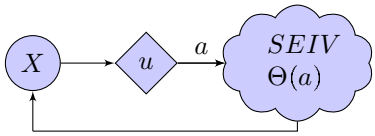


Fig. 2: A diagram of the control architecture studied in this paper. The state X of the process is observed by a controller u, which then applies an action a to the $SEIV$ process, which induces the set of spreading parameters $\Theta(a)$ used to propagate the process forward.

B. Control Architecture

We control the spreading network of the $SEIV$ process by way of applying control actions to the spreading graph of the process. That is, our model assumes the existence of a set of control actions $\mathcal{A}$, about which we assume that each action a fully determines the set of spreading parameters through a known map

$$\Theta(a) = \{\alpha(a), \beta(a), \gamma(a), \delta(a), \eta(a), \xi(a)\},$$

and the cost of applying action a is given by $\mathcal{C}(a)$, where $\mathcal{C}$ is a known function. We study the case in which actions are applied to the network on some predefined set of times $T_{\Delta t} \triangleq \{t \in \mathbb{R}_{\geqslant 0} \,|\, t = \Delta t k, k \in \mathbb{Z}_{\geqslant 0}\}$. The particular action applied is computed by the controller u after observing the state of the system X, where the computation performed is done so as to minimize the economic cost realized by the controller, while still providing some stability guarantee (see Section II-C). Thus, our controller functions as a sampled-data model predictive controller for the continuous-time $SEIV$ process. A diagram presenting this architecture is given in Figure 2.

Whenever the controller decides to change the applied action to a, the parameters of the process dynamics (2) change to $\alpha_i(a)$, $\beta_{ij}(a)$, $\gamma_{ij}(a)$, $\delta_i(a)$, $\eta_i(a)$, and $\xi_i(a)$, respectively, where each value is assumed to be finite, is implicitly defined by Θ, and remains constant between update times. Note that in the remainder of the paper, we will think of the application of a control action as inducing a probability measure which propagates the process forward; for convenience, we notate this measure $\Theta(a)$, as it should cause no confusion.

In an application, the set $\mathcal{A}$ should be chosen so as to represent the types of actions a planner can take in order to effect the evolution of the epidemic, e.g. distributing medication, investing in awareness advertisement, or assigning people to quarantines. The parameter map Θ should be defined to take each possible control action, and produce the particular set of spreading parameters induced by the action. One should

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

expect that giving a person medication should set her healing rate δ_i to a high value, whereas investing in advertisement will set the rate at which a node transitions to the vigilant state (i.e. ξ_i) to a high value, and quarantining a node will result in deleting some connections from the graph (i.e. setting β_{ij} and γ_{ij} to 0) due to persons who ordinarily come into contact with each other not doing so any longer.

We assume that $\mathcal{A}$ is discrete and has at most finitely many elements, and that $\mathcal{C}(a)$ is finite for all $a \in \mathcal{A}$. Note that these assumptions do not play a major role in the text - they do not influence our main results (Theorem 3 and 5) whatsoever. They make the problem setting more concrete, and allow us to define an appropriate optimization method easily, when examining the example application provided in Section V.

Note that there are many works which have studied model predictive control with discrete control actions which precede this text [32]–[37]. As such, this feature should not be considered a primary contribution of the work. In particular, the presence of discrete inputs does not materially change stability analysis significantly in many circumstances, as pointed out in [36]. We believe extending the results in this paper to different action spaces to be straightforward.

C. Economic Model Predictive Control

In contrast to traditional forms of model predictive control, economic model predictive control (EMPC) focuses on developing methods which optimize the economic performance of a system, while still guaranteeing underlying system properties such as stability (see [27], [37]–[46] for relevant background). The controller we study in this paper is an economic model predictive controller which applies feasible, but possibly suboptimal, solutions to the optimization problem

$$\underset{a\in\mathcal{A}}{\text{minimize}}\{\mathcal{C}(a)\,|\,J(X(t),a)\leqslant 0\}, \tag{3}$$

at all times t in the set of sampling times $T_{\Delta t} \triangleq \{t \in \mathbb{R}_{\geqslant 0}\,|\,t = \Delta t k, k \in \mathbb{Z}_{\geqslant 0}\}$. The role of (3) in the closed-loop evolution of the controlled $SEIV$ process is in automatically generating actions from an implicit, nonlinear control law u_{mpc} which limits the rate at which economic costs are incurred by the applied actions while still guaranteeing stability.

Our role in designing the EMPC is in crafting an optimization problem of the form (3) such that applying actions which are feasible to (3) at all times t in $T_{\Delta t}$, the closed-loop behavior of (2) will drive the disease out of the network quickly - say, in an expected amount of time which grows linearly with the size of the initial infection. To this end, we focus on the case in which the *stability constraint function* $J(X(t),a)$ takes the form

$$J(X(t),a) \triangleq \mathbb{E}_{\Theta(a)}\big[\ell(X(t+\Delta t)) - \ell(X(t))e^{-r\Delta t}|X(t)\big], \tag{4}$$

where ℓ is defined as the total number of exposed and infected nodes in the network, i.e. $\ell(X) \triangleq \sum_{i\in\mathcal{V}} X_i^E + X_i^I$, and r is a chosen positive constant describing the desired decay rate of the closed-loop system. The constraint $J(X(t),a) \leqslant 0$, which we refer to as the *stability constraint*, is a common feature of nonlinear and economic model predictive controllers (see [47, Section 3] and [27, Section 3.3], respectively), though the choice of stability constraint function varies depending on the context. Ensuring that the stability constraint is satisfied whenever an action is updated plays a central role in our control by guaranteeing an appropriate notion of decay.

While in full generality, it may be difficult to find a control action which is feasible to (3), we assume that the controller has access to a stabilizing auxiliary control law u_{aux} which provides such an action at any state. Note that this assumption is common in the nonlinear and economic model predictive control literature (see [47, Section 3] and [27, Section 3.3], respectively), due to the abject difficulty of solving (3) to global optimality caused by the nonconvexity induced by the nonlinearity of the dynamics. Moreover, in the context of epidemic containment, we expect finding such controllers to be easy: if we are distributing medication, we may give medication to everyone, if we are distributing protective clothing, we may give protective clothing to everyone, if we are deciding who to quarantine, we may quarantine everyone. In all such cases, if we do so, the epidemic will die out quickly. This is to say that in the context of epidemic containment, we expect that the interesting problem is not generally in finding an abstract control law which will eliminate an epidemic quickly, it is in determining how to apply control actions efficiently, so as to use a nontrivial amount of economic resources while *still* ensuring fast disease elimination.

D. Problem Statement

We address the problem of constructing a robust EMPC for the $SEIV$ process which guarantees that the epidemic attains membership in the set of disease free states quickly, while limiting the rate of resource consumption incurred by the actions applied by the controller. Our work here is focused on addressing two key difficulties which arise from applying economic model predictive control to epidemic containment: (i) approximating the evolution of (2) so as to have a robust approximation of the evolution of the processes' compartmental membership probabilities under a fixed control action a, and (ii) analyzing the convergence of the closed-loop process so as to provide a rigorous guarantee that the disease will be eliminated from the network quickly.

The difficulty of item (i) arises from the fact that in order to induce the desired decay property, the function $J(X(t),a)$ must include information about the future expected state of the process, conditioned on the current state of the process and the action applied, i.e. a term of the form $\mathbb{E}_{\Theta(a)}\big[\ell(X(t+\Delta t))|X(t)\big]$, where ℓ is an appropriately chosen function. Approximating such a function in a manner that guarantees that (3) is feasible requires that we construct a moment closure that can provide a rigorous upper-bound on $\mathbb{E}_{\Theta(a)}\big[\ell(X(t+\Delta t))|X(t)\big]$, which is not a property guaranteed by any prior moment closure techniques. We construct a novel moment closure with this property in Section III.

The difficulty of item (ii) arises from the fact that the stability results which are typical to MPC literature prove asymptotic convergence to a connected, compact set which contains the origin (see, e.g. [26]–[29] and references

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

therein). Such a result is not useful in our context. In general, converging to a set even at an exponential rate does *not* guarantee that a given process *ever* enters the set. As the $SEIV$ process attains membership in the disease-free set in finite time almost surely under all but the most pathological control laws, it is primarily of interest to study how quickly it enters the set of disease free states, which motivates the novel convergence analysis presented in Section IV.

III. Robust Moment Closure for $SEIV$

In this section, we construct a robust moment closure approximation for the $SEIV$ process. Essentially, this is a method for computing an outer approximation to the set of solutions of (2), which we cannot solve directly due to the second-order moments $\mathbb{E}[X_i^C X_j^{C'}]$ not having known, analytic expressions. The method we develop enables rigorously approximating conditional expectations of the form $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$, which are necessary in order to enact the EMPC scheme detailed in Section II-C, as they are required to evaluate the feasibility of (3). Note that here, we only explicitly deal with approximating $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ as defined in Section II-C. This is for brevity, as the techniques can be extended to other choices with more complicated analysis.

The moment closure technique in this section is novel, and has the property that the compartmental membership probabilities of the system are bounded by known quantities for all time, in contrast to prior work. In particular, most prior works in epidemic control literature use a mean-field type moment closure (see, e.g. [31]), which replaces terms of the form $\mathbb{E}[X_i^C X_j^{C'}]$ with the products $\mathbb{E}[X_i^C]\mathbb{E}[X_j^{C'}]$. While such an approximation may work well for sufficiently simple epidemic models [21], in general this approximation gives no rigorous accuracy guarantee, and there are known systems for which such moment closures result in poor approximations of the statistics of the underlying system [22], [48]. Note that nothing formal is known about the quality of mean-field approximations for the $SEIV$ process studied here.

One may imagine that applying more sophisticated types of known moment closure techniques to epidemic processes, such as variants of derivative-matching methods [49]–[52], may avoid this issue. However, this is not the case. Such techniques give only weak guarantees of accuracy with respect to propagating uncertainty forward in time. In particular, the dynamics resulting from such closures are guaranteed to be close to the true dynamics in a neighborhood of the initial condition used to generate the approximation. Such a guarantee is not appropriate for our application, as we must be able to guarantee that our approximation of $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ is an upper-bound of its true value in order to asses the feasibility of (3) to guarantee stability.

A. Fréchet Moment Closure of Moment Dynamics

We begin our formal developments by constructing a naïve robust moment closure for the $SEIV$ expectation dynamics (2). To present the developed system in a concise manner, we first introduce notation for the operators which characterize Fréchet inequalities. These play a central role in our results.

Lemma 1 (Fréchet Inequalities [53]) *Define the operators $\mathcal{F}_\ell(y,z) \triangleq \max\{0, y+z-1\}$, and $\mathcal{F}_u(y,z) \triangleq \min\{y,z\}$. Let $\Pr$ be a probability measure on some event space Ω, and let A and B be events defined on Ω. Then, it holds that*

$$\mathcal{F}_\ell(\Pr(A),\Pr(B)) \leqslant \Pr(A,B) \leqslant \mathcal{F}_u(\Pr(A),\Pr(B)), \quad (5)$$

where the notation $\Pr(A,B)$ denotes the joint probability, i.e. the probability that the event $\{A \cap B\}$ occurs.

The role of the Fréchet inequalities in our development is that of bounding error in the dynamics due to the entanglement caused by the appearance of cross-product terms in (2). Note that the Fréchet inequalities are the tightest inequalities estimating joint probabilities which are functions of only marginal probabilities, and are distribution-free [53]. Note also that our indicator random variables X_i^C are Bernoulli random variables, and so the expectation of their product $\mathbb{E}[X_i^C X_j^{C'}]$ can be written as the joint probability $\Pr(X_i^C = 1, X_j^{C'} = 1)$, making the Fréchet inequalities an appropriate tool for our later analysis. As we cannot analytically propagate the distribution of (1) forward in time, and our dynamics (2) only explicitly give us information about marginal probabilities, this is the best we can get for any given collection of marginals $\{\mathbb{E}[X_i^C]\}_{i\in\mathcal{V},C\in\mathcal{L}}$.

To arrive at a crude approximating system, we use the Fréchet bounds to over-approximate the dynamics on the states we designate as upper bounds, and under-approximate the dynamics on the states we designate as lower bounds. The resulting approximating dynamics, along with their approximation guarantee is given in the following result.

Theorem 1 (Fréchet Moment Closure for $SEIV$) *Let $\underset{\sim}{x}(0) = X(0) = \tilde{x}(0)$, and consider the solutions of the system of nonlinear ordinary differential equations*

$$\begin{aligned}
\dot{\tilde{x}}_i^S =& \alpha_i \tilde{x}_i^V - \xi_i \tilde{x}_i^S - \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_\ell(\tilde{x}_i^S, \underset{\sim}{x}_j^E) - \gamma_{ij}\mathcal{F}_\ell(\tilde{x}_i^S, \underset{\sim}{x}_j^I),\\
\dot{\underset{\sim}{x}}_i^S =& \alpha_i \underset{\sim}{x}_i^V - \xi_i \underset{\sim}{x}_i^S - \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_u(\underset{\sim}{x}_i^S, \tilde{x}_j^E) - \gamma_{ij}\mathcal{F}_u(\underset{\sim}{x}_i^S, \tilde{x}_j^I),\\
\dot{\tilde{x}}_i^E =& \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_u(\tilde{x}_i^S, \tilde{x}_j^E) + \gamma_{ij}\mathcal{F}_u(\tilde{x}_i^S, \tilde{x}_j^I) - \delta_i \tilde{x}_i^E,\\
\dot{\underset{\sim}{x}}_i^E =& \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_\ell(\underset{\sim}{x}_i^S, \underset{\sim}{x}_j^E) + \gamma_{ij}\mathcal{F}_\ell(\underset{\sim}{x}_i^S, \underset{\sim}{x}_j^I) - \delta_i \underset{\sim}{x}_i^E,\\
\dot{\tilde{x}}_i^I =& \delta_i \tilde{x}_i^E - \eta_i \tilde{x}_i^I,\\
\dot{\underset{\sim}{x}}_i^I =& \delta_i \underset{\sim}{x}_i^E - \eta_i \underset{\sim}{x}_i^I,\\
\dot{\tilde{x}}_i^V =& \eta_i \tilde{x}_i^I + \xi_i \tilde{x}_i^S - \alpha_i \tilde{x}_i^V,\\
\dot{\underset{\sim}{x}}_i^V =& \eta_i \underset{\sim}{x}_i^I + \xi_i \underset{\sim}{x}_i^S - \alpha_i \underset{\sim}{x}_i^V.
\end{aligned} \quad (6)$$

Then, for every compartment $C \in \mathcal{L}$ and each node $i \in \mathcal{V}$, the inclusion $\mathbb{E}[X_i^C(t)|X(0)] \in \left[\underset{\sim}{x}_i^C(t|0), \tilde{x}_i^C(t|0)\right]$, holds for all $t \geqslant 0$, where $\underset{\sim}{x}_i^C(t|0)$ and $\tilde{x}_i^C(t|0)$ are used to denote the under- and over-approximation of $\mathbb{E}[X_i^C(t)|X(0)]$ with respect to the dynamics (6), respectively.

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

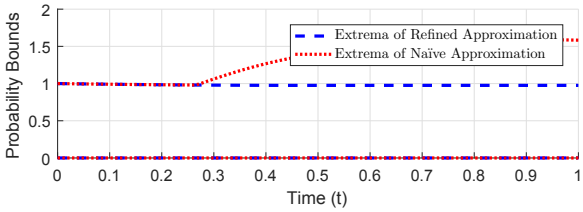


Fig. 3: A demonstration that the crude Fréchet approximation system given by (34) does not generate approximations such that the estimates of compartmental membership probabilities remain bounded in the unit interval, whereas the refined approximation dynamics (10) does. This underscores the fact that when designing a moment closure approximation, care must be taken to ensure that the solutions of the resulting approximation system behave reasonably.

Proof : We state and prove an extension to the comparison lemma (see, e.g., [54, Lemma 3.4]) in Appendix B, which demonstrates a condition under which component-wise orderings are preserved under integration of the dynamics. This is the central technical feature of the proof. Having established an appropriate version of the comparison lemma, all that remains to prove the theorem is to prove that the dynamics (6) satisfy the hypothesis of Lemma 2, where the approximation dynamics (6) are to be compared against the exact dynamics (2). We show this in Appendix C, completing the proof. □

While the approximation dynamics given by (6) are technically correct and *can* be used to estimate $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ rigorously, they are not without fault. In particular, we may note that they include no mechanism for ensuring that the approximations remain bounded in the unit interval. This means that there is no way of ruling out the possibility that the solutions of the approximating dynamics (6) become trivial at some time. Indeed, as demonstrated in Figure 3, it is the case the approximations generated from this system give upper-bounds which exceed one eventually, and hence become trivial. This occurs because the error which enters the approximating system is integrated through time into the evolution of the approximations of the marginal probabilities. Addressing this issue so as to guarantee that the approximations always remain nontrivial requires a deeper analysis, which comes in the following subsection.

B. Refining the Robust Fréchet Moment Closure

We begin this subsection by introducing a proposition which equips us with a formal test for determining when one set of approximation dynamics is better than another, in terms of set inclusion. The result is stated as follows:

Proposition 1 (Constructing Tighter Approximations) *Consider the functions $\underline{f}$ and $\bar{f}$ defining the dynamics (6), and let $\{\underline{x}, \bar{x}\}$ be the solutions of (6). Suppose $\underline{f}$ and $\bar{f}$ are Lipschitz continuous functions on $\mathbb{R}^{(\mathcal{V}\times\mathcal{L})\times 2}$ which satisfy*

$$[\underline{f}_i^C(\underline{x},\bar{x}), \bar{f}_i^C(\underline{x},\bar{x})] \subseteq [\underline{f}_i^C(\underline{x},\tilde{x}), \tilde{f}_i^C(\underline{x},\tilde{x})] \tag{7}$$

on $\bar{\mathcal{X}}_i^C \cup \underline{\mathcal{X}}_i^C$ for all $(i,C) \in \mathcal{V}\times\mathcal{L}$, where $\bar{\mathcal{X}}_i^C$ is the subset of points $(\underline{x},\bar{x},\underline{x},\tilde{x})$ of $\mathbb{R}^{(\mathcal{V}\times\mathcal{L})\times 4}$ such that $\bar{x}_i^C = \tilde{x}_i^C$, and $[\underline{x}_j^{C'}, \bar{x}_j^{C'}] \subseteq [\underline{x}_j^{C'}, \tilde{x}_j^{C'}]$ holds for all $(j,C') \in \mathcal{V}\times\mathcal{L}$, and $\underline{\mathcal{X}}_i^C$ is defined similarly. Then, the system

$$\begin{aligned}\dot{\bar{x}} &= \bar{f}(\underline{x},\bar{x}),\\ \dot{\underline{x}} &= \underline{f}(\underline{x},\bar{x}),\end{aligned} \tag{8}$$

with initial conditions $\underline{x}(0) = \underline{\tilde{x}}(0)$, $\bar{x}(0) = \tilde{x}(0)$ has solutions $\{\underline{x},\bar{x}\}$ which satisfy $[\underline{x}_i^C(t), \bar{x}_i^C(t)] \subseteq [\underline{x}_i^C(t), \tilde{x}_i^C(t)]$ for all $(i,C) \in \mathcal{V}\times\mathcal{L}$ and any time $t \geqslant 0$.

Proof : See Appendix D. □

Proposition 1 gives a means for testing whether or not the crude approximation given by (6) is improved upon by a new candidate approximation, and is the main result we use in developing a refinement of (6).

The most significant modification over (6) required to ensure all estimates remain bounded to the unit interval for all time is in ensuring the dynamics of each approximating upper-bound are non-increasing when the approximating upper-bound is equal to one, and the dynamics of the approximating lower-bound are non-decreasing when it is equal to zero. To accomplish such an approximation, we introduce complement bounds to the dynamics, which we define formally as follows.

Definition 1 (Complement Bounding Operator) Let $C \in \{S,E,I,V\}$ be a compartmental label of the $SEIV$ spreading process. We define the complement upper-bounding operator associated to C as the nonlinear operator

$$\bar{\mathcal{B}}_i^C y \triangleq \min\{1-\bar{x}_i^C, y\}. \tag{9}$$

In essence, we see that improvements over (6) can be made if we replace all instances of variables which can cause unbounded growth with appropriate complement bounds. By doing so systematically, we arrive at a better approximating system, which we give here in Theorem 2.

Theorem 2 (Refined Fréchet Moment Closure of $SEIV$) *Let $\underline{x}(0) = X(0) = \bar{x}(0)$, and consider the solutions of the system of nonlinear ordinary differential equations*

$$\begin{aligned}
\dot{\bar{x}}_i^S &= \alpha_i \bar{\mathcal{B}}_i^S \bar{x}_i^V - \xi_i \bar{x}_i^S - \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_\ell(\bar{x}_i^S, \underline{x}_j^E) - \gamma_{ij}\mathcal{F}_\ell(\bar{x}_i^S, \underline{x}_j^I),\\
\dot{\underline{x}}_i^S &= \alpha_i \underline{x}_i^V - \xi_i \underline{x}_i^S - \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_u(\underline{x}_i^S, \bar{x}_j^E) - \gamma_{ij}\mathcal{F}_u(\underline{x}_i^S, \bar{x}_j^I),\\
\dot{\bar{x}}_i^E &= \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_u(\bar{\mathcal{B}}_i^E \bar{x}_i^S, \bar{x}_j^E) + \gamma_{ij}\mathcal{F}_u(\bar{\mathcal{B}}_i^E \bar{x}_i^S, \bar{x}_j^I) - \delta_i \bar{x}_i^E,\\
\dot{\underline{x}}_i^E &= \sum_{j\in\mathcal{N}_i} \beta_{ij}\mathcal{F}_\ell(\underline{x}_i^S, \underline{x}_j^E) + \gamma_{ij}\mathcal{F}_\ell(\underline{x}_i^S, \underline{x}_j^I) - \delta_i \underline{x}_i^E,\\
\dot{\bar{x}}_i^I &= \delta_i \bar{\mathcal{B}}_i^I \bar{x}_i^E - \eta_i \bar{x}_i^I,\\
\dot{\underline{x}}_i^I &= \delta_i \underline{x}_i^E - \eta_i \underline{x}_i^I,\\
\dot{\bar{x}}_i^V &= \eta_i \bar{\mathcal{B}}_i^V \bar{x}_i^I + \xi_i \bar{\mathcal{B}}_i^V \bar{x}_i^S - \alpha_i \bar{x}_i^V,\\
\dot{\underline{x}}_i^V &= \eta_i \underline{x}_i^I + \xi_i \underline{x}_i^S - \alpha_i \underline{x}_i^V
\end{aligned} \tag{10}$$

Then, for every compartment $C \in \mathcal{L}$ and each node $i \in \mathcal{V}$,

$$\mathbb{E}[X_i^C(t)|X(0)] \in \left[\underline{x}_i^C(t|0), \bar{x}_i^C(t|0)\right] \subseteq [0,1], \tag{11}$$

and

$$\left[\underline{x}_i^C(t|0), \bar{x}_i^C(t|0)\right] \subseteq \left[\underline{x}_i^C(t|0), \tilde{x}_i^C(t|0)\right], \tag{12}$$

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

hold for all $t \geqslant 0$, *where* $\underline{x}_i^C(t|0)$ *and* $\bar{x}_i^C(t|0)$ *are used to denote the under- and over-approximation of* $\mathbb{E}[X_i^C(t)|X(0)]$ *with respect to the dynamics* (10), *and* $\underline{x}_i^C(t|0)$ *and* $\tilde{x}_i^C(t|0)$ *are used to denote the under- and over-approximation of* $\mathbb{E}[X_i^C(t)|X(0)]$ *with respect to the dynamics* (6).

Proof : See Appendix E. □

We can think of the process that we have used to arrive at (10) from (2) as one of successively pruning the set of trajectories permitted by the approximating systems. In the step where we moved from (2) to (6), we used the Fréchet probability bounds to constrain the set of trajectories our system may admit as solutions to a superset of the set of probability measures the Fréchet approximations permit. In moving from (6) to (10), we further restrict the set of solutions to those which satisfy simple complementarity bounds. In so doing, each step improved the accuracy with which the dynamics are approximated. We now consider how to use (10) in order to approximate the term $E_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ in the constraint function.

C. Decay Constraint Approximation

It is obvious that if we set $\underline{x}(0) = X(0) = \bar{x}(0)$ and integrate (10) over the interval $[0, t]$, we get that

$$\mathbb{E}_{\Theta(a)}[\ell(X(t))|X(0)] \leqslant \ell(\bar{x}(t)) \tag{13}$$

holds. However, (13) is not the best approximation that can be derived from the solutions of (10). The optimal approximation which can be obtained from propagating the dynamics (10) can be found efficiently via linear programming, as we state formally in the following result and its proof.

Theorem 3 (Optimal Approximation of Decay Function) *Let* $\underline{x}(0) = X(0) = \bar{x}(0)$, *and consider the solutions* $(\underline{x}, \bar{x})$ *of* (10) *evaluated at time* t. *It holds that*

$$\mathbb{E}_{\Theta(a)}[\ell(X(t))|X(0)] \leqslant \psi(X(0), a) \triangleq \sum_{i\in\mathcal{V}} \min\{\bar{x}_i^E(t|0) + \bar{x}_i^I(t|0), 1 - \underline{x}_i^S(t|0) - \underline{x}_i^V(t|0)\}, \tag{14}$$

where the bound is the tightest which can be derived from the inclusions generated by integrating (10), *and in particular is pointwise tighter than the bound* (13).

Proof : See Appendix F. □

From Theorem 3, we can see that while the states of the dynamical system (10) are not themselves the best approximation possible for the conditional expectation $\mathbb{E}_{\Theta(a)}[\ell(X(t))|X(0)]$, the optimal approximation can be defined as a nonlinear output function $\sum_{i\in\mathcal{V}} \min\{\bar{x}_i^E(t|0) + \bar{x}_i^I(t|0), 1 - \underline{x}_i^S(t|0) - \underline{x}_i^V(t|0)\}$, and can thus be computed as efficiently as integrating the dynamics (10). Since we may repeat the arguments used in support of Theorem 3 just as easily with the initial time taking the value t, it is clear that we can efficiently compute an upper-bound on $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ from any state, at any time. We use this approximation scheme in a simulated application in Section V, where we see (Figure 4) that the approximation generates nontrivial approximations that are useful for controlling the $SEIV$ process.

IV. Convergence Analysis

In this section, we analyze the convergence of the controller induced by applying suboptimal solutions of (3), where we choose $J(X(t), a)$ as in (4), and $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ is approximated by the method described in Theorem 3 (see Section III-C). Abstractly, we are most concerned with ensuring that the epidemic attains membership in the set of disease-free states quickly. To formalize this notion, we define the *elimination time* of the process as follows:

Definition 2 (Elimination Time) The elimination time τ_{elim} of an $SEIV$ epidemic is the first time at which all nodes are in neither the exposed nor the infected compartment, i.e. $\tau_{\text{elim}} \triangleq \inf\{t \geqslant 0 \,|\, \ell(X(t)) = 0\}$, where $\ell(X(t)) \triangleq \sum_{i\in\mathcal{V}} X_i^E(t) + X_i^I(t)$.

We seek a guarantee on the expected elimination time of the $SEIV$ process under the designed EMPC scheme. Since the $SEIV$ process, as well as most compartmental epidemic processes in general, attains the disease-free set in finite time almost surely under all but the most pathological control laws, we must consider a strong notion of stability for it to be meaningful. Here, we show that the expectation τ_{elim} grows slowly with the size of the initial infection.

The analysis that we perform to arrive at an upper bound for the expectation of the elimination time τ_{elim} relies critically on knowledge of the evolution of the expected number of exposed and infected nodes in the graph. As such, we first analyze this expectation (in Section IV-A), and then analyze the expectation of the elimination time (in Section IV-B).

A. Bounding the Expected Number of Exposed and Infected Nodes

We find that the proposed EMPC method uniformly exponentially eliminates the epidemic in expectation, with respect to the set of sampling times $T_{\Delta t}$, stated formally as follows.

Theorem 4 (Exponential Elimination Under EMPC) *Let* $X(0)$ *be the initial state of the* $SEIV$ *process, choose any* $r > 0$, *and any* $\Delta t > 0$. *Suppose an auxiliary control policy* u_{aux} *exists such that*

$$\mathbb{E}_{u_{\text{aux}}}[\ell(X(t+\Delta t))|X(t)] \leqslant \ell(X(t))e^{-r\Delta t}, \tag{15}$$

for all $t \in T_{\Delta t}$. *Then, the evolution of the* $SEIV$ *process under the policy* u_{mpc} *generated by the economic model predictive controller specified in Section II-C satisfies*

$$\mathbb{E}_{u_{mpc}}[\ell(X(t))|X(0)] \leqslant \ell(X(0))e^{-rt}, \tag{16}$$

for all $t \in T_{\Delta t}$. *Moreover, the bound* (16) *is tight.*

Proof : See Appendix G. □

The proof of Theorem 4 follows from an induction argument, which makes appeals to the expectation decay constraint (15), fundamental tools from the theory of probability, and the decay property encoded in the actions contained in feasible set of (3) through the constraint $J(X(t), a) \leqslant 0$. With respect to our problem, the principle importance of Theorem 4 is in allowing us to rigorously analyze the elimination time of the process, which we perform in the following subsection.

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

B. Bounding the Expected Elimination Time

While it is intuitive that an $SEIV$ process in which the total count of exposed and infected nodes decays exponentially quickly might have a small elimination time, there is no immediately apparent link between the two concepts. Moreover, the decay property guaranteed by Theorem 4 is *not* uniform exponential elimination; the exponential decay is only guaranteed on a countable subset of times. As such, we may only infer anything about the expected number of exposed and infected nodes on a small subset of times, and must use this information to prove a result on the elimination time of the process. We know of no previously published technique for doing so in the literature, but fortunately it can be done, as shown in the following result and its proof:

Theorem 5 (Bound on Expected Elimination Time) *Suppose the $SEIV$ process evolves under the policy u_{mpc} generated by the EMPC method detailed in Section II-C, where $J(X(t),a)$ is defined by* (4)*, and $\mathbb{E}_{\Theta(a)}[\ell(X(t+\Delta t))|X(t)]$ is approximated by the method detailed in Theorem 3. Then, the expected elimination time satisfies*

$$\mathbb{E}_{u_{\text{mpc}}}[\tau_{elim}|X(0)] \leqslant \tau_1 + \frac{\mathrm{e}^{-r\tau_1}}{1-\mathrm{e}^{-r\Delta t}}\Delta t \ell(X(0)), \qquad (17)$$

where τ_1 is the first time in the sampling time set such that the expected number of exposed and infected nodes is less than one, which can be shown to be

$$\tau_1 = \left\lceil \frac{\log(\ell(X(0)))}{r\Delta t} \right\rceil \Delta t, \qquad (18)$$

where $\lceil a \rceil$ denotes the smallest integer larger than a, i.e. the ceil of a. Moreover, the bound (17) *is tight.*

Proof : The essence of this argument relies on approximating the integral which defines the expected elimination time. The approximation occurs in three steps:(i) representing the expectation as an integral of the distribution function of the elimination time random variable, (ii) finding a convergent, closed-form approximation to the distribution function of the elimination time random variable, and (iii) evaluating the integral of the approximated distribution function.

Since τ_{elim} is a non-negative random variable which takes values on the real line, we may use the layer-cake representation of expectation to identify the equivalence

$$\mathbb{E}_{u_{\text{mpc}}}[\tau_{\text{elim}}|X(0)] \triangleq \int_0^{\infty} 1 - F_{\tau_{\text{elim}}}(h)dh \qquad (19)$$

where $F_{\tau_{\text{elim}}}$ is the distribution function of τ_{elim}, i.e. $F_{\tau_{\text{elim}}}(t) \triangleq \Pr(\tau_{\text{elim}} \leqslant t)$. Since once $\ell(X(\tau)) = 0$ for some τ, it holds for all $t \geqslant \tau$, we have the identity $\Pr(\tau_{\text{elim}} \leqslant t) = \Pr(\ell(X(t)) = 0)$, which is equivalent to the expression $1 - \Pr(\tau_{\text{elim}} \leqslant t) = \Pr(\ell(X(t)) > 0)$.

As such, we may construct an upper-bound on $1 - F_{\tau_{\text{elim}}}(t)$ by constructing an upper-bound on $\Pr(\ell(X(t)) > 0)$. We do so using the decay properties already proven of the designed controller. In particular, Theorem 4 gives that

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X(t))|X(0)] \leqslant \ell(X(0))e^{-rt} \qquad (20)$$

holds for all $t \in T_{\Delta t}$, and so we may upper-bound $\Pr(\ell(X(t)) > 0)$ for all $t \in T_{\Delta t}$ by the optimal value of

$$\underset{p \in \Psi_{[n]_0}}{\text{maximize}} \quad \sum_{k=1}^{n} p_k \qquad (21a)$$

$$\text{subject to} \quad \sum_{k=0}^{n} k p_k \leqslant \ell(X(0))e^{-r\left\lfloor \frac{t}{\Delta t} \right\rfloor \Delta t}, \qquad (21b)$$

where $\lfloor a \rfloor$ denotes the largest integer smaller than a, and $\Psi_{[n]_0}$ is the set of all possible marginal probability assignments over $[n]_0$, so chosen because the number of exposed and infected nodes in the graph must take value on $[n]_0$, and the constraint (21b) enforces the expectation inequality (20). Since the left-hand-side of inequality (21b) is least sensitive to increases in p_1 for any value p_k with $k \geqslant 1$, one may show that the optimal value of (21) can be computed as $\min\{1, \ell(X(0))e^{-r\left\lfloor \frac{t}{\Delta t} \right\rfloor \Delta t}\}$. Hence,

$$\Pr(\ell(X(t)) > 0 | X(0)) \leqslant \min\{1, \ell(X(0))\mathrm{e}^{-r\left\lfloor \frac{t}{\Delta t} \right\rfloor \Delta t}\} \quad (22)$$

holds for all $t \in T_{\Delta t}$, where the bound (22) is tight.

Since the distribution function $F_{\tau_{\text{elim}}}$ is non-decreasing with respect to t by definition, and $1 - F_{\tau_{\text{elim}}}(t) = \Pr(\ell(X(t)) > 0)$, we have that $\Pr(\ell(X(t)) > 0)$ is *non-increasing* with respect to t. Hence (22) holds for *all* times $t \geqslant 0$. Thus, (19) and (22) together imply the inequality

$$\mathbb{E}_{u_{\text{mpc}}}[\tau_{\text{elim}}|X(0)] \leqslant \int_0^{\infty} \min\left\{1, \ell(X(0))\mathrm{e}^{-r\left\lfloor \frac{h}{\Delta t} \right\rfloor \Delta t}\right\} dh. \qquad (23)$$

We now seek a closed-form expression of the right hand side of (23). Defining τ_1 as

$$\tau_1 \triangleq \inf\{t \in \mathbb{R}_{\geqslant 0} \,|\, \ell(X(0))\mathrm{e}^{-r\left\lfloor \frac{t}{\Delta t} \right\rfloor \Delta t} \leqslant 1\},$$

we can evaluate τ_1 to satisfy the claimed identity (18), and can then rewrite (23) as

$$\mathbb{E}_{u_{\text{mpc}}}[\tau_{\text{elim}}|X(0)] \leqslant \tau_1 + \sum_{j=\frac{\tau_1}{\Delta t}}^{\infty} \ell(X(0))\Delta t \left(\mathrm{e}^{-r\Delta t}\right)^j. \quad (24)$$

Evaluating the tail of the geometric sum in (24) gives

$$\mathbb{E}_{u_{\text{mpc}}}[\tau_{\text{elim}}|X(0)] \leqslant \tau_1 + \frac{\mathrm{e}^{-r\tau_1}}{1-\mathrm{e}^{-r\Delta t}}\Delta t \ell(X(0)),$$

which is as stated in the theorem's hypothesis.

Finally, note that since the bound on the distribution function $F_{\tau_{\text{elim}}}$ derived is optimal at all times among all such bounds which use only the bound on the expected number of exposed and infected nodes in the graph provided by Theorem 4, and the bound provided by Theorem 4 is itself tight, it follows as well that the upper bound (23) is optimal among all such guarantees that can be provided by the designed EMPC. This completes the proof. □

Since (23) grows at worst linearly with respect to the number of initially infected and exposed nodes in the graph, the derived bound certifies that the designed EMPC scheme eliminates the epidemic from the network quickly. Thus, the problem stated in Section II-D has been appropriately solved.

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

Note that the proof given for Theorem 5 relies critically on the fact that the expected number of infected and exposed nodes decays exponentially quickly, as if this were not the case, the approximation used for the integrand would not be integrable, and the resulting approximation would be trivial. Moreover, the argument relies on the topology of the state-space of epidemic processes in order to guarantee that the controller attains membership in the targeted set of states quickly. Indeed, if it were not the case that the optimal value of (21) is bounded away from one after only a short amount of time, our attempt at approximating the expectation of τ_{elim} meaningfully would fail as well. Since this occurs *only* because $\ell(X)$ *must* take values on the set $[n]_0$, it follows that such a convergence argument will not generalize to EMPC schemes on general state spaces, but *will* generalize to other epidemic process readily, as all such processes taking place on finite graphs evolve on finite state spaces.

Remark 1 (Use of Alternate Approximation Methods) As noted in Section III, the motivation for constructing a robust moment closure as we have is to be certain that the convergence guarantees we arrive at give us information on the behavior of the statistics of the $SEIV$ process. By adequately accounting for the worst-case introduction of approximation errors in the dynamics as we have, this goal was accomplished. However, one may wish to use alternate approximations of the dynamics in regions of the state space wherein robustness is unimportant, in order to attempt to find actions which improve cost performance.

For example, if for an n node graph, we partition $\mathcal{X}$ into sets $\mathcal{W}_k = \{Y \in \mathcal{X} \,|\, \ell(Y) = k\}$ for each $k \in [n]_0$, we may decide to use a mean-field type moment closure or a Monte Carlo simulation to approximate $J(X(t), a)$ when $X \in \cup_{k=0}^{\tilde{k}} \mathcal{W}_k$ for some $\tilde{k} < n$, and use the robust approximations derived in Section III otherwise. In this situation, the controller developed here will drive X into $\cup_{k=0}^{\tilde{k}} \mathcal{W}_k$ quickly whenever it leaves the set. That is, such a controller will be guaranteed to keep the total number of exposed and infected nodes below $\tilde{k}$ efficiently, and will otherwise use other approximations in order to attempt to use fewer resources. The math required to formalize and prove the above claim is very similar to that which was used to prove the convergence guarantees in this section, and as such we will not present it here. •

V. Application: Optimizing Quarantine Use

In this section, we present a concrete application for the developed EMPC framework. Note that while we only present one application here, the general principles contained in Sections III and IV are not constrained to this context. Indeed, whenever anyone should want to consider a new application, all one needs to do is specify an action space $\mathcal{A}$, a parameter map Θ, an auxiliary control law u_{aux}, and an appropriate optimization method.

A. Quarantine Model for $SEIV$

We consider the problem of strategically removing nodes from the spreading graph in order to efficiently drive an $SEIV$ epidemic to extinction quickly. This is a mathematical model for the practical problem of deciding who to quarantine, and for how long, in the presence of an epidemic contagion. Note that in the context of this problem, the maximum realized resource use rate is the number of quarantine beds needed throughout the course of the epidemic, and as such provides a reasonable index for evaluating the cost of the controller. Furthermore, given that world governments are currently in the process of providing disease control agencies with sweeping authority to quarantine individuals exposed to infectious disease (see, e.g., the recent U.S. bill [55]) despite recommendations from the medical community (see, e.g., [56]), understanding the mathematics of when quarantining is necessary for the control of a disease is of utmost importance.

We represent control actions here by an n-dimensional vector a in which $a_i = 1$ if and only if the i'th node is removed from the spreading graph (i.e., quarantined), and $a_i = 0$ otherwise. With this notation, we may represent our action space as $\mathcal{A} = \{0, 1\}^n$. We model quarantining a node by removing its outgoing edges from the spreading graph, i.e. we have for all pairs (i, j), the exposure rates have the functional forms $\beta_{ij}(a_j) \triangleq \bar{\beta}_{ij} - \bar{\beta}_{ij} a_j$, and $\gamma_{ij}(a_j) \triangleq \bar{\gamma}_{ij} - \bar{\gamma}_{ij} a_j$, where each a_j is restricted to the set $\{0, 1\}$. For simplicity, we assume that cost of quarantining nodes is additive, and so may be represented as $\mathcal{C}(a) \triangleq \sum_{i \in \mathcal{V}} a_i$. In this context, the value of the cost function evaluated for a particular control action a is representative of the number of beds required to implement the quarantine strategy. Applying economic model predictive control to this problem explicitly attempts to minimize the number of beds used in execution.

Note that our work here is *not* the first to study the problem of quarantine management for models of biological disease. Typical works from the pre-existing literature study quarantine management problems for mean-field epidemic models, and model a node being in quarantine by adding an additional compartment to the compartmental spreading model, which does not interact with any other compartments (see, e.g., [18], [19] for specific instances). The control design is done by way of varying the rate at which nodes transition to the quarantine compartment, with actuating the rate coming at a given cost. Our model here is similar, in that our quarantined nodes do not interact with the rest of the network, and placing the node in quarantine comes at a cost to the controller. Note, however, that our work here is the first work which considers quarantine optimization for stochastic networked epidemics, and in this sense is novel.

B. An Auxiliary Control Law for Quarantine Optimization

As noted in Section II-C, our control scheme assumes the existence of an auxiliary control law which always satisfies the required expectation decay constraint. While specifying an all-purpose auxiliary control law is outside of the scope of this paper, we demonstrate here how to construct one for the quarantining problem, with the hope that it will provide insight on how to do so in other application areas.

The policy we design for the quarantine problem, which we refer to as the total quarantine policy, removes nodes

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

which are either exposed to or infected by the disease at each time that the state of the process is observed. Intuitively, this is a mathematical model for what is implemented in the event of a serious disease outbreak (e.g., the response to the Ebola epidemic of 2014 [56]). This procedure is guaranteed mathematically to eliminate the contagion from the network exponentially quickly provided the control horizon is sufficiently long, as we show in the following result.

Theorem 6 (Convergence of Total Quarantine Policy) *Suppose η_i and δ_i are distinct for all $i \in \mathcal{V}$, choose $r < \min\{\eta_i, \delta_i\}$ for all i, and Δt to satisfy*

$$\frac{\log(\max\{\delta_i, \eta_i\}) - \log(|\eta_i - \delta_i|)}{\min\{\delta_i, \eta_i\} - r} \leqslant \Delta t, \tag{25}$$

for all $i \in \mathcal{V}$. Suppose further that at each $t \in T_{\Delta t}$, an action from the total quarantine policy u_{tot},

$$u_{tot}(X) \triangleq \begin{cases} a_i = 1, & \forall i \text{ s.t. } X_i^E + X_i^I > 0, \\ a_i = 0, & \forall i \text{ s.t. } X_i^E + X_i^I = 0, \end{cases} \tag{26}$$

is applied and held constant until time $t + \Delta t$. Then, the evolution of the $SEIV$ process satisfies

$$\mathbb{E}_{u_{tot}}[\ell(X(t + \Delta t))|X(t)] \leqslant \ell(X(t))\mathrm{e}^{-r\Delta t} \tag{27}$$

for all $X(t) \in \mathcal{X}$, and each $t \in T_{\Delta t}$.

Proof : See Appendix H. □

Note that the statement requiring η_i and δ_i to be distinct is one of explanatory convenience. In particular, our analysis relies on solving a particular system of linear ordinary differential equations, of which η_i and δ_i are the eigenvalues. By requiring that they be distinct, we simplify the required proofs to only having to consider one possible type of solution. Of course, the analysis can be done just as easily in the case that $\delta_i = \eta_i$. However, the principle components of the argument are the same, and thus are left out of the paper. It should also be noted that the sampling time bound given by (25) is conservative. This inequality was derived for the express purpose of providing a simple inequality which can be checked easily, at the expense of applying approximations.

The total quarantine policy analyzed in Theorem 6 is conservative. It removes more nodes from the network than is required to eliminate the epidemic exponentially quickly. The numerical experiments in Section V-D verify that the proposed EMPC finds actions which are more efficient than the total quarantine policy.

C. A Method for Optimizing Quarantine Use

We now consider the task of of finding good approximate solutions to (3). Note that since the $SEIV$ process evolves on a state space with 4^n elements, evaluating the stability constraint $J(X, a) \leqslant 0$ precisely is in general difficult. As such, we use the results developed in Section III to evaluate the constraint conservatively. In particular, at a state $X(0)$ and for a particular control action a, define the function

$$\bar{J}(X(0), a) = \psi(X(0), a) - \ell(X(0))e^{-r\Delta t}, \tag{28}$$

where the term $\psi(X(0), a)$ is defined as in Theorem 3. Since $\bar{J}(X(0), a) \geqslant J(X(0), a)$, it follows that $\bar{J}(X(0), a) \leqslant 0$ implies $J(X(0), a) \leqslant 0$. Hence, computing an approximate solution which is feasible to

$$\underset{a \in \mathcal{A}}{\text{minimize}}\{\mathcal{C}(a) \,|\, \bar{J}(X(t), a) \leqslant 0\}, \tag{29}$$

also computes an approximate solution which is feasible to (3). Note that (29) is itself still not easy to solve: it is an integer programming problem. While there are relatively sophisticated methods for solving various classes of integer programs, they all require the presence of some special structure to work well, as integer programs are in general $\mathcal{NP}$-hard. Moreover, even computing suboptimality bounds for arbitrary integer programming problem is difficult [57], [58]. As such, we use a randomized multistart search to approximately solve (29). This is detailed in Algorithm 1.

Algorithm 1 Multi-Start Local Descent

Initialization:

Define feasible auxiliary solution $a_{\text{aux}} = u_{\text{aux}}(X(t))$;
Initialize set of candidate solutions $\mathbb{A}^\star \triangleq \{a_{\text{aux}}\}$;
Define maximum iteration count $k_{\max}$;
Set $k = 0$.
Run main program;

Main Program:

while $k \leqslant k_{\max}$ **do**
  Sample a from distribution with support $\mathcal{A}$;
  if a is feasible **then**
    $\mathcal{I} \triangleq \{i \in \mathcal{V} \,|\, a_i = 1\}$;
    if $a - e_i$ is feasible for some $i \in \mathcal{I}$ **then**
      $a \leftarrow a - e_i$;
      Go to 3;
    else
      $\mathbb{A}^\star \leftarrow \mathbb{A}^\star \cup \{a\}$.
    end if
  else
    Return a infeasible;
  end if
  $k \leftarrow k + 1$;
end while
Return $a^\star = \operatorname{argmin}_{a \in \mathbb{A}^\star}\{\mathcal{C}(a)\}$;

While in general it may take a very long time to find the optimal solution, it is theoretically guaranteed that Algorithm 1 will find the globally optimal solution of (3) eventually, and in finite time, so long as $k_{\max}$ is set to infinity. This is moreso a nice theoretical guarantee, than a practically important matter. Of course, the amount of time required to find a feasible solution of a particular quality will be affected by the particular choice of sampling distribution, and it is likely that the optimal solution will not be found for quite some time. We detail the particular sampling distribution used in our experiments in Section V-D.

Perhaps most important in practice is that for any considered candidate action, Algorithm 1 will terminate at a

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

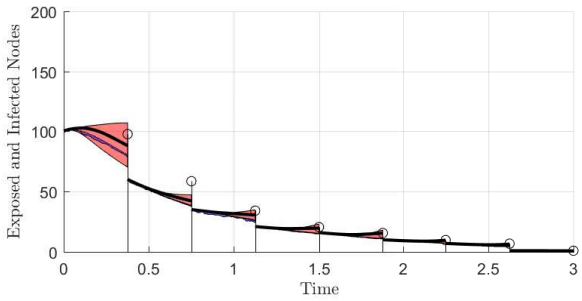
locally optimal point after at most $O(n^2)$ operations, which can be proven formally by a simple counting argument. This guarantees that at any given system state, our controller can do better than random guessing quickly, while not necessarily guaranteeing that the optimal action will be found. Note that by no means is this the only sort of algorithm which can be used here. Rather, this is the simplest approach which has been found to work well enough to be worth reporting here. It is expected that in different application domains, researchers may want to investigate the efficacy of different approaches. Those interested may want to read up on heuristic algorithms for integer programming problems (see, e.g., [59], [60]).

D. Numerical Experiments

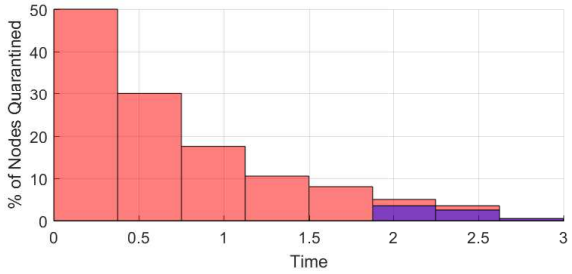
We simulate the evolution of the $SEIV$ process under the EMPC defined in Section II-C, using the optimization method defined in Section V-C with decay rate $r = 0.07$ and $\Delta t = 0.375$, where the candidate solutions for Algorithm 1 are chosen such that exposed and infected nodes are quarantined independently with probability 0.7, and susceptible and vigilant nodes are quarantined independently with probability 0.1. All numerical integrations are performed with Matlab's implementation of `ode45`. We compare its performance against that of the total quarantine base heuristic given in Section V-B. Representative results of our numerical study are given in Figure 4, which reports the result of a simulation of a 200-node Erdös-Reyni random graph with connection probability 0.6, and spreading parameters chosen as $\alpha_i = 0.1$, $\beta_{ij} = 0.1$, $\gamma_{ij} = 0.1$, $\delta_i = 1.25$, $\eta_i = 3.5$, and $\xi_i = 2$, where the parameters are equivalued for all edges and nodes so as to be able to fully specify the problem considered.

From the convergence plot given in Figure 4a, we see that while the approximations generated by the solutions of (10) are somewhat loose at the beginning of the simulation, they are highly nontrivial by themselves. Moreover, the optimal upper-bound given by Theorem 3 improves the approximation further, decreasing the uncertainty of the approximation by a few units in some cases. At their worst, the optimal bounds constrain the expectation of the process to within an interval of approximately 20 nodes, corresponding to an uncertainty of approximately 10%, with respect to the number of nodes in the graph. While this suggests that an ideal controller may be able to attain better performance, it also suggests that possible improvements are limited. As time passes, the quality of the approximation improves to the point where the upper- and lower- approximations converge, further exemplifying the approximation method's utility.

From the cost plot given in Figure 4b, we can see that the cost of the controller is substantially reduced, with respect to a comparison against the total quarantine base policy. To some extent, this provides a mathematical validation of the opinions expressed by the medical community, which suggest that quarantining individuals exposed to infectious diseases is *not always* required for effective disease control [56]. Of course, the $SEIV$ model we study here is perhaps too simple to say anything more concrete.



(a) Total number of exposed and infected nodes in the spreading graph under the proposed controller as a function of time. The shaded red regions give the bounds generated by propagating the dynamics (10), the open circles indicate the upper-bound on the expected number of exposed and infected nodes guaranteed by the controller, the solid black line gives the optimal upper-bound as computed by Theorem 3, the shaded blue regions give the 98% confidence intervals generated from estimating the expectation of the process by producing $10k$ sample paths of the process via Monte Carlo simulation, and then computing $1k$ estimates of the mean via bootstrap sampling from the simulated trajectories.



(b) Fraction of nodes held in quarantine as a function of time. The red regions indicate the total fraction of nodes quarantined under the total quarantine policy, the blue regions indicate the total fraction of nodes quarantined under the proposed controller. In this particular simulation, a quarantine is not required at all until the number of infections in the graph is very small.

Fig. 4: Plots numerically evaluating the performance of the proposed controller. Figure 4a shows that the proposed controller induces exponential elimination of the epidemic. Figure 4b demonstrates that the stochastic optimization method proposed in Algorithm 1 can significantly reduce the number of nodes required to be quarantine in order to guarantee the desired elimination rate.

VI. Summary and Future Work

In this paper, we have developed a robust economic model predictive controller for the mitigation of diseases modeled by $SEIV$ processes. In addressing this problem, we provided a novel robust moment closure technique, which guarantees that the engendered approximations always give rigorous, nontrivial upper- and lower- bounds of the processes' compartmental membership probabilities. We have also shown how to analyze the convergence properties of the proposed EMPC in order to guarantee that the elimination time of the process is appropriately small. In so doing, we have provided a useful step in the process to developing control techniques for continuous-time stochastic networked epidemics which provide statistical guarantees about the evolution of the process. Still, there is much work to be done.

Perhaps an obvious criticism of the basic $SEIV$ model is that - due to modeling the underlying contact and compartment transition processes as Poisson processes - the

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

holding time distributions for each node's compartmental memberships are constricted to be exponential, which may not be reflective of what is seen in actual diseases. As such, it would be of particular interest to extend the analysis of the process presented in this paper to a non-Markovian setting, such that non-exponential holding time distributions can be accommodated. It is also of interest to determine which types of optimization methods work best for controllers designed for applications in different domains, where the objectives and epidemic models may be different. In particular, it seems of interest to extend the framework presented here to the types of problems which have been studied exclusively in the mean-field regime. This paper provides a step in the direction of solving such problems; we believe that continued interest from the community will lead to interesting solutions.

APPENDIX

A. Derivation of $SEIV$ Expectation Dynamics

Since the integrands of the Itô integrals of (1) are always finite, the process is square integrable. As such, a consequence of Itô's lemma is that the expectation operator and the Itô integral commute (see, e.g., [61, Theorem 3.20]), and the expectation of the probability measures of the Poisson processes become the rates of the process. Carrying this computation through, we may take the expectation of both sides of (1) to arrive at the integral equations

$$\begin{aligned}
\mathbb{E}[X_i^S(t)] &= \int_0^t \mathbb{E}[X_i^V(h)]\alpha_i - \mathbb{E}[X_i^S(h)]\xi_i \\
&\quad - \sum_{j\in\mathcal{N}_i} (\mathbb{E}[X_i^S X_j^E(h)]\beta_{ij} + \mathbb{E}[X_i^S X_j^I(h)]\gamma_{ij})\mathrm{d}h, \\
\mathbb{E}[X_i^E(t)] &= \int_0^t \sum_{j\in\mathcal{N}_i} (\mathbb{E}[X_i^S X_j^E(h)]\beta_{ij} + \mathbb{E}[X_i^S X_j^I(h)]\gamma_{ij}) \\
&\quad - \mathbb{E}[X_i^E(h)]\delta_i \mathrm{d}h, \\
\mathbb{E}[X_i^I(t)] &= \int_0^t \mathbb{E}[X_i^E(h)]\delta_i - \mathbb{E}[X_i^I(h)]\eta_i \mathrm{d}h, \\
\mathbb{E}[X_i^V(t)] &= \int_0^t \mathbb{E}[X_i^I(h)]\eta_i + \mathbb{E}[X_i^S(h)]\xi_i \\
&\quad - \mathbb{E}[X_i^V(h)]\alpha_i \mathrm{d}h.
\end{aligned} \tag{30}$$

Note that only ordinary Riemann integrals remain in (30), so we may apply the fundamental theorem of calculus (see, e.g., [62, Theorem 6.20]) to arrive at the system of ordinary, nonlinear differential equations (2). □

B. Two-Sided Multivariate Comparison Lemma

Lemma 2 (Two-Sided Multivariate Comparison Lemma) *Consider a system of differential equations*

$$\dot{x} = f(x) \tag{31}$$

with $x \in \mathbb{R}^p$, $f : \mathbb{R}^p \to \mathbb{R}^p$, and possessing a unique, continuously differentiable solution $x(t)$. Suppose $\tilde{f}$ and $\underset{\sim}{f}$ are Lipschitz continuous vector functions defined on $\mathbb{R}^{p\times 2}$, where for each component i,

$$f_i(z) \leqslant \tilde{f}_i(\underset{\sim}{z}, \tilde{z}) \tag{32}$$

holds everywhere on the set

$$\tilde{\mathcal{Z}}_i \triangleq \{(\underset{\sim}{z}, z, \tilde{z}) \in \mathbb{R}^{(p\times 3)} \,|\, \underset{\sim}{z}_i \leqslant z_i = \tilde{z}_i, \underset{\sim}{z}_j \leqslant z_j \leqslant \tilde{z}_j, \forall j \neq i\},$$

and the inequality

$$\underset{\sim}{f}_i(\underset{\sim}{z}, \tilde{z}) \leqslant f_i(z) \tag{33}$$

holds everywhere on the set

$$\underset{\sim}{\mathcal{Z}}_i \triangleq \{(\underset{\sim}{z}, z, \tilde{z}) \in \mathbb{R}^{(p\times 3)} \,|\, \underset{\sim}{z}_i = z_i \leqslant \tilde{z}_i, \underset{\sim}{z}_j \leqslant z_j \leqslant \tilde{z}_j, \forall j \neq i\}.$$

Then, the solutions to the system

$$\begin{aligned}
\dot{\tilde{x}} &= \tilde{f}(\underset{\sim}{x}, \tilde{x}), \\
\dot{\underset{\sim}{x}} &= \underset{\sim}{f}(\underset{\sim}{x}, \tilde{x}),
\end{aligned} \tag{34}$$

with initial conditions $\underset{\sim}{x}(0) = x(0) = \tilde{x}(0)$, satisfy

$$x_i(t) \in [\underset{\sim}{x}_i(t), \tilde{x}_i(t)] \tag{35}$$

for all i and $t \geqslant 0$.

Proof : The proof proceeds by a sequence of contradiction arguments, which in particular use the fact that the solutions $\underset{\sim}{x}(t)$, $x(t)$, and $\tilde{x}(t)$ are continuously differentiable along with the inequalities (33) and (32) in order to demonstrate that the inclusions $x_i(t) \in [\underset{\sim}{x}_i(t), \tilde{x}_i(t)]$ hold for all $i \in [p]$, and all times $t \geqslant 0$. Note that the continuous differentiability of the solutions follows immediately from Lipschitz continuity of the dynamics (see, e.g., [54]).

Suppose for purposes of contradiction that $\underset{\sim}{x}(t) \leqslant x(t) \leqslant \tilde{x}(t)$ does not hold for all time. That then implies that there is some t at which either $x_i(t) > \tilde{x}_i(t)$ or $x_i(t) < \underset{\sim}{x}_i(t)$ occurs. Let t_c be the first time at which such an event occurs. Suppose for now that $x_i(t_\mathrm{c}) > \tilde{x}_i(t_\mathrm{c})$ occurs; we argue the other case analogously. By the continuity and differentiability of x_i and $\tilde{x}_i$, it then also holds that there exists some time t_c^- such that $t_\mathrm{c}^- < t_\mathrm{c}$, $x_i(t_\mathrm{c}^-) = \tilde{x}_i(t_\mathrm{c}^-)$, and $\dot{x}_i(t_\mathrm{c}^-) > \dot{\tilde{x}}_i(t_\mathrm{c}^-)$. However, since it holds that $\underset{\sim}{x}_i(t_\mathrm{c}^-) \leqslant x_i(t_\mathrm{c}^-) = \tilde{x}_i(t_\mathrm{c}^-)$, and we have that $\underset{\sim}{x}_j(t_\mathrm{c}^-) \leqslant x_j(t_\mathrm{c}^-) \leqslant \tilde{x}_j(t_\mathrm{c}^-)$, we have by assumption that $\dot{x}_i(t_\mathrm{c}^-) \leqslant \dot{\tilde{x}}_i(t_\mathrm{c}^-)$. This is a contradiction. The case in which $x_i(t_\mathrm{c}) < \underset{\sim}{x}_i(t_\mathrm{c})$ occurs can be handled by similar arguments. This completes the proof. □

C. Proof that the dynamics (6) satisfy Lemma 2

Note that by using a construction such as [31, Section 5], the $SEIV$ process can be represented as a 4^n-dimensional time-homogeneous Markov process for fixed set of spreading parameters. As such, the compartmental membership probabilities generated from any particular initial state X can be represented as sums of states evolving as solutions to a 4^n-dimensional linear system, and so are unique and continuously differentiable. Likewise, since each term of the dynamics (6) is a sum of Lipschitz continuous functions, it follows that the dynamics (6) are also Lipschitz continuous. It remains to verify that the inequalities required by Lemma 2 are satisfied on the appropriate subsets of the state space.

Choose some node label $i \in \mathcal{V}$ and some compartmental label $C \in \mathcal{L}$. We wish to show that

$$\max_{\xi\in\Xi(x)} \left\{ \frac{\mathrm{d}\mathbb{E}_\xi[X_i^C]}{\mathrm{dt}} \right\} \leqslant \dot{\tilde{x}}_i^C(\underset{\sim}{x}, \tilde{x}) \tag{36}$$

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

holds everywhere on

$$\tilde{\mathcal{X}}_i^C \triangleq \{(\underline{x}, x, \tilde{x}) \in \mathbb{R}^{(\mathcal{V}\times\mathcal{L})\times 3} \mid \underline{x}_i^C \leqslant x_i^C = \tilde{x}_i^C, \underline{x}_j^{C'} \leqslant x_j^{C'} \leqslant \tilde{x}_j^{C'}, (j, C') \in \mathcal{V}\times\mathcal{L}\},$$

where the set $\Xi(x)$ is the set of all probability measures with first moments $x_i^C = \mathbb{E}_\xi[X_i^C]$, and where the term $\dot{\tilde{x}}_i^C(\underline{x}, \tilde{x})$ is a shorthand reference to the dynamics (6). Consider all terms of the function $\frac{\mathrm{d}\mathbb{E}_\xi[X_i^C]}{\mathrm{dt}}$ with positive coefficients; they can be written as $\kappa\mathbb{E}_\xi[X_i^C X_j^{C'}]$ for some $\kappa \geqslant 0$. Each is bounded above by $\mathcal{F}_u(\tilde{x}_i^C, \tilde{x}_j^{C'})$ by the Fréchet inequality (5) for any measure with expectation x. Likewise, consider all terms of the function $\frac{\mathrm{d}\mathbb{E}_\xi[X_i^C]}{\mathrm{dt}}$ with negative coefficients; they can be written as $\kappa\mathbb{E}_\xi[X_i^C X_j^{C'}]$ for some $\kappa \leqslant 0$. From the Fréchet inequality (5), we have $\kappa\mathcal{F}_\ell(\tilde{x}_i^C, \underline{x}_j^{C'}) \geqslant \kappa\mathbb{E}_\xi[X_i^C X_j^{C'}]$, so long as $\tilde{x}_i^C = \mathbb{E}_\xi[X_i^C]$, which is precisely the case for points on $\tilde{\mathcal{X}}_i^C$. Hence, the inequality (36) holds on $\tilde{\mathcal{X}}_i^C$, as claimed. We can show that the inequality

$$\dot{\underline{x}}_i^C(\underline{x}, \tilde{x}) \leqslant \min_{\xi\in\Xi(x)} \left\{ \frac{\mathrm{d}\mathbb{E}_\xi[X_i^C]}{\mathrm{dt}} \right\}$$

holds everywhere on the set

$$\underline{\mathcal{X}}_i^C \triangleq \{(\underline{x}, x, \tilde{x}) \in \mathbb{R}^{(\mathcal{V}\times\mathcal{L})\times 3} \mid \underline{x}_i^C = x_i^C \leqslant \tilde{x}_i^C, \underline{x}_j^{C'} \leqslant x_j^{C'} \leqslant \tilde{x}_j^{C'}, (j, C') \in \mathcal{V}\times\mathcal{L}\}$$

by similar arguments, which completes the proof. □

D. Proof of Proposition 1

This result relies directly on a well-known comparison lemma from the theory of monotone dynamical systems, the Kamke-Müller lemma (see, e.g. [63]). In its most basic form, the Kamke-Müller lemma gives comparisons between the solutions of one dynamical system, evolving from two distinct initial conditions which satisfy some ordering. However, it is well known that this type of comparison can be made for the solutions of two distinct dynamical systems, provided certain ordering conditions hold. Specifically, the result we use in this paper is stated as follows, where we have adapted material from [63, Section 3.1] to the notation used here:

Lemma 3 (Extended Kamke-Müller Lemma) *Let f and g be Lipschitz continuous vector functions on $\mathbb{R}^p$, and consider the nonlinear dynamical system*

$$\dot{y} = f(y)$$
$$\dot{z} = g(z).$$

If for each i, the inequality $f_i(\underline{q}) \leqslant g_i(\bar{q})$ holds on

$$\mathcal{Q}_i \triangleq \left\{(\underline{q}, \bar{q}) \in \mathbb{R}^{p\times 2} \mid \underline{q}_i = \bar{q}_i, \underline{q}_j \leqslant \bar{q}_j, j \neq i\right\},$$

and it holds that $y(t_0) = z(t_0)$, then it also holds that $y_i(t) \leqslant z_i(t)$ for all $t \geqslant t_0$, and all $i \in [p]$.

Now, consider taking $y^T = [\bar{x}^T, -\underline{x}^T]^T$, and $z^T = [\tilde{x}^T, -\underline{x}^T]^T$, and let $p = 4n$, where we let each $j \in [4n]$ represent exactly one $(i, C) \in \mathcal{V}\times\mathcal{L}$. If for each $(i, C) \in \mathcal{V}\times\mathcal{L}$, we have that $[\underline{f}_i^C(\underline{q}), \bar{f}_i^C(\underline{q})] \subseteq [\underline{f}_i^C(\bar{q}), \tilde{f}_i^C(\bar{q})]$, on the corresponding set $\mathcal{Q}_j$, and $(\underline{x}(0), \bar{x}(0)) = (\underline{x}(0), \tilde{x}(0))$, Lemma 3 implies $[\underline{x}_i^C(t), \bar{x}_i^C(t)] \subseteq [\underline{x}_i^C(t), \tilde{x}_i^C(t)]$ for all $(i, C) \in \mathcal{V}\times\mathcal{L}$, and for all $t \geqslant 0$, as claimed. □

E. Proof of Theorem 2

It is easy to check that for all nodes i, and all compartmental labels $C \in \mathcal{L} \triangleq \{S, E, I, V\}$, we have that

$$[\underline{f}_i^C(\underline{x}, \bar{x}), \bar{f}_i^C(\underline{x}, \bar{x})] \subseteq [\underline{f}_i^C(\underline{x}, \tilde{x}), \tilde{f}_i^C(\underline{x}, \tilde{x})],$$

holds everywhere on $\bar{\mathcal{X}}_i^C \cup \underline{\mathcal{X}}_i^C$ as defined in Proposition 1, which then gives that $[\underline{x}_i^C(t), \bar{x}_i^C(t)] \subseteq [\underline{x}_i^C(t), \tilde{x}_i^C(t)]$ holds for any initial condition $\underline{x}(0) = \tilde{x}(0) = \underline{x}(0) = \bar{x}(0) = x(0)$, all $t \geqslant 0$, and all $(i, C) \in \mathcal{V}\times\mathcal{L}$. Since including the complement bounding operators still gives valid over- and under-approximations of the terms of $\frac{\mathrm{d}\mathbb{E}[X_i^C]}{\mathrm{dt}}$ on the relevant subsets of the state space, this certifies that the system (10) generates solutions which satisfy $\mathbb{E}[X_i^C(t)] \in [\underline{x}_i^C(t), \bar{x}_i^C(t)]$ for all $t \geqslant 0$, and all $(i, C) \in \mathcal{V}\times\mathcal{L}$.

To demonstrate that $[\underline{x}(t), \bar{x}(t)] \subseteq [0, 1]^{\mathcal{V}\times\mathcal{L}}$ holds for all t, we make an appeal to the continuity of the solutions $(\underline{x}(t), \bar{x}(t))$ as well the values of $\dot{\underline{x}}$ and $\dot{\bar{x}}$ on the sets of states $(\underline{x}, \bar{x})$, such that $\underline{x}_i^C = 0$ or $\bar{x}_i^C = 1$ for some i or C. Suppose for purposes of contradiction that $\underline{x}_i^C(t_\mathrm{c}) < 0$ for some time t_c. By continuity of $\underline{x}_i^C$, it must then be the case that there exists some time $t_\mathrm{c}^- < t_\mathrm{c}$ such that $\underline{x}_i^C(t_\mathrm{c}^-) = 0$ and $\dot{\underline{x}}_i^C(t_\mathrm{c}^-) < 0$. However, evaluating the expression for $\dot{\underline{x}}_i^C$ with $\underline{x}_i^C = 0$, has $\dot{\underline{x}}_i^C \geqslant 0$. This is a contradiction, and as such, proves that $\underline{x}_i^C(t) \geqslant 0$ for all $t \geqslant 0$.

Suppose now that $\bar{x}_i^C(t_\mathrm{c}) > 1$ for some finite time t_c. By continuity of $\bar{x}_i^C(t_\mathrm{c})$, it must then be the case that there exists some time $t_\mathrm{c}^- < t_\mathrm{c}$ such that $\bar{x}_i^C(t_\mathrm{c}^-) = 1$ and $\dot{\bar{x}}_i^C(t_\mathrm{c}^-) > 0$. However, evaluating the expression for $\dot{\bar{x}}_i^C$ with $\bar{x}_i^C = 1$, has $\dot{\bar{x}}_i^C \leqslant 0$. This is a contradiction, and as such, proves that $\bar{x}_i^C(t) \leqslant 1$ for all $t \geqslant 0$. Since each component of the dynamics (10) is a sum of Lipschitz continuous functions, they are Lipschitz continuous. This completes the proof. □

F. Proof of Theorem 3

Let $\Psi_\mathcal{G}$ be the set of all possible marginal compartmental membership probabilities for the graph, i.e.

$$\Psi_\mathcal{G} \triangleq \left\{ y \in \mathbb{R}_{\geqslant 0}^{\mathcal{V}\times\mathcal{L}} \mid \sum_{C\in\mathcal{L}} y_i^C = 1, \forall\, i \in \mathcal{V} \right\} \tag{37}$$

Define Λ_t as the set of marginal compartmental membership probabilities permitted by the approximations generated by integrating the dynamics (10) with initial condition $\underline{x}(0) = X(0) = \bar{x}(0)$ over the interval $[0, t]$, i.e.

$$\Lambda_t \triangleq \left\{ y \in \Psi_\mathcal{G} | y_i^C \in [\underline{x}_i^C(t), \bar{x}_i^C(t)], \forall i \in \mathcal{V}, C \in \mathcal{L} \right\}.$$

Since $\ell(X)$ is a sum of indicator random variables, it follows that $\mathbb{E}[\ell(X(t))|X(0)]$ can take the value $\ell(y)$ for any $y \in \Lambda_t$. By maximizing over all such y, it follows that $\mathbb{E}[\ell(X(t))|X(0)] \leqslant \max_{y\in\Lambda_t}\{\ell(y)\}$ holds, and is tight because $y \in \Lambda_t$. It remains to show that $\max_{y\in\Lambda_t}\{\ell(y)\}$ evaluates to the right hand side of (14).

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

To do this, we note that $\max_{y\in\Lambda_t}\{\ell(y)\}$ can be decomposed as the sum of the n linear programs

$$\begin{aligned} \underset{y_i\in\mathbb{R}^{\mathcal{L}}_{\geqslant 0}}{\text{maximize}} \quad & y_i^E + y_i^I && \text{(38a)}\\ \text{subject to} \quad & y_i^C \in [\underline{x}_i^C, \bar{x}_i^C], \forall i \in \mathcal{V}, C \in \mathcal{L} && \text{(38b)}\\ & \sum_{C\in\mathcal{L}} y_i^C = 1 && \text{(38c)} \end{aligned}$$

We now solve (38) analytically. By considering the constraint (38c), we have that at all feasible points of (38) satisfy $y_i^E + y_i^I = 1 - y_i^S - y_i^V$. As such, if $\bar{x}_i^E + \bar{x}_i^I$ and $1 - \underline{x}_i^V + \underline{x}_i^S$ take distinct values, only the smaller of the two values is attainable on the feasible polytope. As the objective is monotonically increasing in y_i^E and y_i^I, it then follows that the optimal value of (38) is bounded above by $\min\{\bar{x}_i^E + \bar{x}_i^I, 1 - \underline{x}_i^V - \underline{x}_i^S\}$. Noting that the program (38) is guaranteed to be feasible since the interval constraints (38b) are generated so as to contain the true underlying marginal probabilities, we can finish the argument by showing that there exists a feasible point which attains the value $\min\{\bar{x}_i^E + \bar{x}_i^I, 1 - \underline{x}_i^V - \underline{x}_i^S\}$.

Suppose that $\bar{x}_i^E + \bar{x}_i^I > 1 - \underline{x}_i^V - \underline{x}_i^S$. It follows that $\bar{x}_i^E + \bar{x}_i^I + \underline{x}_i^V + \underline{x}_i^S > 1$, from which feasibility of (38) and the fact that y_i^S and y_i^V are lower bounded by $\underline{x}_i^S$ and $\underline{x}_i^V$ (respectively) implies that there exists a y_i such that $y_i^E + y_i^I = 1 - \underline{x}_i^V + \underline{x}_i^S$. This point attains the value $\min\{\bar{x}_i^E + \bar{x}_i^I, 1 - \underline{x}_i^V - \underline{x}_i^S\}$. The case in which $\bar{x}_i^E + \bar{x}_i^I \leqslant 1 - \underline{x}_i^V - \underline{x}_i^S$ can be handled by similar arguments. This ends the proof. □

G. Proof of Theorem 4

Our argument proceeds by demonstrating that the sequence of controls realized by the proposed EMPC method induce the desired decay property. Principally, the proof relies on an induction, and an application of the tower property of conditional expectations [64, Proposition 13.2.7].

We have by construction that the inequality

$$\mathbb{E}_{u_{\text{aux}}}[\ell(X(t+\Delta t))|X(t)] \leqslant \ell(X(t))\mathrm{e}^{-r\Delta t} \qquad (39)$$

holds. Since the optimization routine will only pass back solutions which satisfy the stability constraint, as it will at worst pass back the action taken by the auxiliary policy, we have that

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X(t+\Delta t))|X(t)] \leqslant \ell(X(t))\mathrm{e}^{-r\Delta t} \qquad (40)$$

holds as well. We use this in an induction to prove the inequality demanded by the theorem's statement. Take

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X(\Delta t))|X(0)] \leqslant \ell(X(0))\mathrm{e}^{-r\Delta t} \qquad (41)$$

as a base for induction, and for an induction hypothesis that

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X(k\Delta t))|X(0)] \leqslant \ell(X(0))\mathrm{e}^{-rk\Delta t} \qquad (42)$$

holds for some arbitrary positive integer k. We now show that this implies that

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X((k+1)\Delta t))|X(0)] \leqslant \ell(X(0))\mathrm{e}^{-r(k+1)\Delta t}.$$

By expanding the conditioning in accordance with the tower property of conditional expectation [64, Proposition 13.2.7], we have that the identity

$$\begin{aligned} &\mathbb{E}_{u_{\text{mpc}}}[\ell(X((k+1)\Delta t))|X(0)]\\ &\quad = \mathbb{E}_{u_{\text{mpc}}}[\mathbb{E}[\ell(X((k+1)\Delta t))|X(k\Delta t), X(0)] \end{aligned} \qquad (43)$$

holds. From definition, we have the inequality

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X((k+1)\Delta t))|X(k\Delta t), X(0)] \leqslant \ell(X(k\Delta t))\mathrm{e}^{-r\Delta t},$$

which when applied to the identify (43) gives

$$\begin{aligned} &\mathbb{E}_{u_{\text{mpc}}}[\mathbb{E}[\ell(X((k+1)\Delta t))|X(k\Delta t), X(0)]]\\ &\qquad \leqslant \mathbb{E}_{u_{\text{mpc}}}[\ell(X(k\Delta t))|X(0)]\mathrm{e}^{-r\Delta t}. \end{aligned}$$

Now, the induction hypothesis (42) yields the inequality

$$\mathbb{E}_{u_{\text{mpc}}}[\ell(X((k+1)\Delta t))|X(0)]\mathrm{e}^{-r\Delta t} \leqslant \ell(X(0))\mathrm{e}^{-r(k+1)\Delta t},$$

which shows that (16) holds for any $t \in T_{\Delta t}$. Noting that this final inequality is tight in the case where (15) is tight at every sampling time completes the proof. □

H. Proof of Theorem 6

The essence of this proof is in demonstrating that the total quarantine policy induces sufficient negative drift in the process so as to guarantee the expectation decay stated in the theorem's hypothesis. To accomplish this, we analyze the evolution of the upper-bounds on the compartmental membership probabilities for node i and compartments E and I. Writing their dynamics down from (6) with β_{ij} and γ_{ij} set to 0 for all j, we have that the linear system

$$\begin{aligned} \dot{\bar{x}}_i^E &= -\delta_i \bar{x}_i^E,\\ \dot{\bar{x}}_i^I &= \delta_i \bar{x}_i^E - \eta_i \bar{x}_i^I, \end{aligned} \qquad (44)$$

describes the behavior of the approximating system under the total quarantine policy. Using standard solution techniques from the theory of linear ordinary differential equations to solve (44) with initial conditions $\bar{x}^E(t)$ and $\bar{x}^I(t)$ and η_i, δ_i distinct gives the solutions:

$$\begin{aligned} &\bar{x}_i^I(\Delta t|t) + \bar{x}_i^E(\Delta t|t) =\\ &\bar{x}_i^I(t)\mathrm{e}^{-\eta_i\Delta t} + \frac{\eta_i}{\eta_i - \delta_i}\bar{x}_i^E(t)\mathrm{e}^{-\delta_i\Delta t} - \frac{\delta_i}{\eta_i - \delta_i}\bar{x}_i^E(t)\mathrm{e}^{-\eta_i\Delta t}. \end{aligned}$$

Because we observe the state $X(t)$ at each sampling time $t \in T_{\Delta t}$, we have only two possible initial conditions for each $i : \bar{x}_i^E(t) = 1$ and $\bar{x}_i^I(t) = 0$ or $\bar{x}_i^E(t) = 0$ and $\bar{x}_i^I(t) = 1$.

In the first case, we need to verify that there exists some $r > 0$ and Δt which satisfy the decay constraint

$$\frac{\eta_i}{\eta_i - \delta_i}\mathrm{e}^{-\delta_i\Delta t} - \frac{\delta_i}{\eta_i - \delta_i}\mathrm{e}^{-\eta_i\Delta t} \leqslant \mathrm{e}^{-r\Delta t}. \qquad (45)$$

By assuming $\eta_i > \delta_i > r$ rearranging terms, approximating the resulting inequality, and taking logarithms, we get

$$\frac{\ln(\eta_i) - \ln(\eta_i - \delta_i)}{\delta_i - r} \leqslant \Delta t. \qquad (46)$$

Similarly, by assuming $\delta_i > \eta_i > r$, we get the inequality

$$\frac{\ln(\delta_i) - \ln(\delta_i - \eta_i)}{\eta_i - r} \leqslant \Delta t. \qquad (47)$$

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

Considering both inequalities together verifies that the inequality stated by the theorem's hypothesis, i.e.

$$\frac{\ln(\max\{\eta_i,\delta_i\}) - \ln(|\eta_i - \delta_i|)}{\min\{\eta_i,\delta_i\} - r} \leqslant \Delta t \tag{48}$$

suffices to demonstrate that when (44) is initialized with $\bar{x}_i^E(t) = 0$ and $\bar{x}_i^I(t) = 1$, we satisfy the desired exponential decay inequality. In the case where $\bar{x}_i^E(t) = 1$ and $\bar{x}_i^I(t) = 0$, we have that the exponential decay inequality is satisfied for any t, and all $r < \eta_i$, which is implied by the inequality claimed by the hypothesis.

Recollecting our argument, we see that for r and Δt satisfying (48), we have the decay inequality

$$\bar{x}_i^E(\Delta t|t) + \bar{x}_i^I(\Delta t|t) \leqslant (\bar{x}_i^E(t) + \bar{x}_i^I(t))\mathrm{e}^{-r\Delta t}.$$

By summing over all of the nodes in the network, we get

$$\sum_{i=1}^{n} \bar{x}_i^E(\Delta t|t) + \bar{x}_i^I(\Delta t|t) \leqslant \left(\sum_{i=1}^{n} \bar{x}_i^I(t) + \bar{x}_i^E(t)\right) \mathrm{e}^{-r\Delta t}$$

which since $\ell(X(t)) = \left(\sum_{i=1}^n \bar{x}_i^I(t) + \bar{x}_i^E(t)\right)$ and $\mathbb{E}[\ell(X(t+\Delta t))|X(t)] \leqslant \left(\sum_{i=1}^n \bar{x}_i^I(\Delta t|t) + \bar{x}_i^E(\Delta t|t)\right)$ hold, together imply

$$\mathbb{E}_{u_{\text{tot}}}[\ell(X(t+\Delta t))|X(t)] \leqslant \ell(X(t))\mathrm{e}^{-r\Delta t}$$

which was sought. This concludes the proof. □

REFERENCES

[1] N. T. J. Bailey, *The Mathematical Theory of Infectious Diseases and Its Applications*. High Wycombe, U.K.: Charles Griffin & Company, Ltd., 2nd ed., 1975.

[2] V. M. Preciado, M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas, "Optimal Resource Allocation for Network Protection Against Spreading Processes," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 99–108, 2014.

[3] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Optimal Resource Allocation for Control of Networked Epidemic Models," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 2, pp. 159–169, 2017.

[4] M. Ogura and V. M. Preciado, "Stability of Spreading Processes over Time-Varying Large-Scale Networks," *IEEE Transactions on Network Science and Engineering*, vol. 3, no. 1, pp. 44–57, 2016.

[5] S. Wang, M. H. R. Khouzani, B. Krishnamachari, and F. Bai, "Optimal Control for Epidemic Routing of Two Files with Different Priorities in Delay Tolerant Networks," in *IEEE American Control Conference*, (Chicago, IL, USA), pp. 1387–1392, 2015.

[6] P. Lee, A. Clark, B. Alomair, L. Bushnell, and R. Poovendran, "Adaptive Mitigation of Multi-Virus Propagation: A Passivity-Based Approach," vol. 5870, no. c, pp. 1–11, 2016.

[7] N. J. Watkins, C. Nowzari, V. M. Preciado, and G. J. Pappas, "Optimal resource allocation for competing epidemics over arbitrary networks," in *Proceedings of the American Control Conference*, pp. 1381–1386, 2015.

[8] S. Eshghi, M. H. Khouzani, S. Sarkar, and S. S. Venkatesh, "Optimal Patching in Clustered Malware Epidemics," *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 283–298, 2016.

[9] S. Eshghi, S. Sarkar, and S. S. Venkatesh, "Visibility-Aware Optimal Contagion of Malware Epidemics," *IEEE Transactions on Automatic Control*, vol. 62, no. October, pp. 5205–5212, 2017.

[10] A. R. Hota and S. Sundaram, "Optimal network topologies for mitigating security and epidemic risks," *54th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2016*, pp. 1129–1136, 2017.

[11] P. E. Paré, J. Liu, C. L. Beck, A. Nedić, and T. Başar, "Multi-Competitive Viruses over Static and Time – Varying Networks," pp. 1–16.

[12] K. Drakopoulos, A. Ozdaglar, and J. Tsitsiklis, "An Efficient Curing Policy for Epidemics on Graphs," *IEEE Transactions on Network Science and Engineering*, vol. 1, no. 2, pp. 67–75, 2015.

[13] K. Scaman, A. Kalogeratos, and N. Vayatis, "Suppressing Epidemics in Networks Using Priority Planning," *IEEE Transactions on Network Science and Engineering*, vol. 3, no. 4, pp. 271–285, 2016.

[14] N. J. Watkins, C. Nowzari, and G. J. Pappas, "Inference, Prediction, and Control of Networked Epidemics," in *IEEE American Control Conference*, (Seattle, WA, USA), pp. 5611 – 5616, 2017.

[15] X. Wei, N. C. Valler, B. Aditya Prakash, I. Neamtiu, M. Faloutsos, and C. Faloutsos, "Competing Memes Propagation on Networks: A Network Science Perspective," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 6, pp. 1049–1060, 2013.

[16] M. J. Keeling and K. T. Eames, "Networks and epidemic models," *Journal of The Royal Society Interface*, vol. 2, no. 4, pp. 295–307, 2005.

[17] C. Eksin, J. S. Shamma, and J. S. Weitz, "Disease dynamics in a stochastic network game: a little empathy goes a long way in averting outbreaks," *Scientific Reports*, vol. 7, no. July 2016, p. 44122, 2017.

[18] F. Brauer, "Some Simple Epidemic Models," *Mathematical Biosciences and Engineering*, vol. 3, no. 1, pp. 1–15, 2006.

[19] X. Yan, Y. Zou, and J. Li, "Optimal quarantine and isolation strategies in epidemics control," *World Journal of Modelling and Simulation*, vol. 3, no. 3, pp. 202–211, 2007.

[20] C. Nowzari, V. M. Preciado, and G. J. Pappas, "Analysis and Control of Epidemics," *IEEE Control Systems Magazine*, vol. 36, no. 1, pp. 26–46, 2016.

[21] P. Simon and I. Z. Kiss, "On bounding exact models of epidemic spread on networks," *arXiv preprint*, pp. 1–18, 2017.

[22] N. J. Watkins, C. Nowzari, V. M. Preciado, and G. J. Pappas, "Optimal Resource Allocation for Competitive Spreading Processes on Bilayer Networks," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 1, pp. 298–307, 2018.

[23] T. K. Dasaklis, C. P. Pappis, and N. P. Rachaniotis, "Epidemics control and logistics operations: A review," *International Journal of Production Economics*, vol. 139, no. 2, pp. 398–410, 2012.

[24] C. J. E. Metcalf, R. B. Birger, S. Funk, R. D. Kouyos, J. O. Lloyd-Smith, and V. A. A. Jansen, "Five challenges in evolution and infectious diseases," *Epidemics*, vol. 10, pp. 40–44, 2015.

[25] C. A. Arias and B. E. Murray, "Antibiotic-Resistant Bugs in the 21st Century — A Clinical Super-Challenge," *The New England Journal of Medicine*, pp. 439–443, 2009.

[26] F. Allgöwer and A. Zheng, eds., *Nonlinear model predictive control*. Boston: Springer Basel AG, 2000.

[27] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

[28] A. Mesbah, "Stochastic Model Predictive Control," *IEEE Control Systems Magazine*, vol. 3, no. December, pp. 30–44, 2016.

[29] F. Borelli, A. Bemporad, and M. Morari, *Predictive Control for linear and hybrid systems*. Cambridge University Press, first ed., 2017.

[30] M. M. Hikal, "Dynamic Properties for a General SEIV Epidemic," *Electronic Journal of Mathematical Analysis and Applications*, vol. 2, no. 1, pp. 26–36, 2014.

[31] F. D. Sahneh, C. Scoglio, and P. V. Mieghem, "Generalized Epidemic Mean-Field Model for Spreading Processes Over Multilayer Complex Networks," *IEEE/ACM Transactions on Networking*, vol. 21, no. 5, pp. 1609–1620, 2013.

[32] T. Geyer and D. E. Quevedo, "Multistep finite control set model predictive control for power electronics," *IEEE Transactions on Power Electronics*, vol. 29, no. 12, pp. 6836–6846, 2014.

[33] M. Preindl and S. Bolognani, "Model predictive direct torque control with finite control set for pmsm drive systems, part 2: Field weakening operation," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 2, pp. 648–657, 2013.

[34] P. Karamanakos, T. Geyer, N. Oikonomou, F. D. Kieferndorf, and S. Manias, "Direct model predictive control: A review of strategies that achieve long prediction intervals for power electronics," *IEEE Industrial Electronics Magazine*, vol. 8, no. 1, pp. 32–43, 2014.

[35] A. Richards and J. How, "Mixed-integer programming for control," *Proceedings of the 2005, American Control Conference, 2005.*, pp. 2676–2683, 2005.

[36] J. B. Rawlings and M. J. Risbeck, "Model predictive control with discrete actuators: Theory and application," *Automatica*, vol. 78, pp. 258–265, 2017.

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TAC.2019.2919136, IEEE Transactions on Automatic Control

[37] J. B. Rawlings, N. R. Patel, M. J. Risbeck, C. T. Maravelias, M. J. Wenzel, and R. D. Turney, "Economic MPC and real-time decision making with application to large-scale HVAC energy systems," *Computers and Chemical Engineering*, vol. 114, pp. 89–98, 2018.

[38] M. Ellis, J. Liu, and P. Christophides, *Economic Model Predictive Control*. Springer, 2016.

[39] D. Angeli, R. Amrit, and J. B. Rawlings, "Enforcing convergence in nonlinear economic MPC," *Proceedings of the IEEE Conference on Decision and Control*, pp. 3387–3391, 2011.

[40] R. Amrit, J. B. Rawlings, and D. Angeli, "Economic optimization using model predictive control with a terminal cost," *Annual Reviews in Control*, vol. 35, no. 2, pp. 178–186, 2011.

[41] K. Subramanian, J. B. Rawlings, and C. T. Maravelias, "Economic model predictive control for inventory management in supply chains," *Computers and Chemical Engineering*, vol. 64, pp. 71–80, 2014.

[42] R. Amrit, J. B. Rawlings, and L. T. Biegler, "Optimizing process economics online using model predictive control," *Computers and Chemical Engineering*, vol. 58, pp. 334–343, 2013.

[43] D. Angeli, R. Amrit, and J. B. Rawlings, "On average performance and stability of economic model predictive control," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1615–1626, 2012.

[44] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control*. Springer, 2017.

[45] T. Faulwasser, L. Grüne, and M. A. Müller, "Economic Nonlinear Model Predictive Control," *Foundations and Trends in Systems and Control*, vol. 5, no. 1, p. 90, 2018.

[46] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017.

[47] D. Q. Mayne, "Nonlinear Model Predictive Control: Challenges and Opportunites," in *Nonlinear Model Predictive Control* (F. Allgöwer and A. Zheng, eds.), ch. 2, pp. 23–44, Boston: Springer Basel AG, first ed., 2000.

[48] E. Cator and P. Van Mieghem, "Second-order mean-field susceptible-infected-susceptible epidemic threshold," *Physical Review E*, vol. 85, no. 5, pp. 1–7, 2012.

[49] A. Singh and J. P. Hespanha, "Moment closure techniques for stochastic models in population biology," *American Control Conference*, pp. 6—-pp, 2006.

[50] A. Singh and J. P. Hespanha, "Approximate moment dynamics for chemically reacting systems," *IEEE Transactions on Automatic Control*, vol. 56, no. 2, pp. 414–418, 2011.

[51] K. R. Ghusinga, M. Soltani, A. Lamperski, S. Dhople, and A. Singh, "Approximate moment dynamics for polynomial and trigonometric stochastic systems," *arXiv preprint*.

[52] M. Soltani, C. Vargas, N. Kumar, R. Kulkarni, and A. Singh, "Moment Closure Approximations in a Genetic Negative Feedback Circuit," p. 6, 2014.

[53] L. Rüschendorf, "Fréchet-Bounds and Their Applications," Tech. Rep. January, 1991.

[54] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, third ed., 2002.

[55] CDC/HHS, "Final Rule for Control of Communicable Diseases," 2017.

[56] J. M. Drazen, R. Kanapathipillai, E. W. Campion, E. J. Rubin, S. M. Hammer, S. Morrissey, and L. R. Baden, "Ebola and Quarantine," *New England Journal of Medicine*, vol. 371, no. 21, pp. 2029–2030, 2014.

[57] S. Burer and A. N. Letchford, "Non-convex mixed-integer nonlinear programming: A survey," *Surveys in Operations Research and Management Science*, vol. 17, no. 2, pp. 97–106, 2012.

[58] J. Lee, Y. Sun, and M. Saunders, "Proximal Newton-type methods for convex optimization," *Nips*, pp. 1–9, 2012.

[59] C. Blum and G. R. Raidl, *Hybrid Metaheuristics: Powerful Tools for Optimization*. Springer, 2016.

[60] P. Hansen and N. Mladenović, "Variable neighborhood search: Principles and applications," *European Journal of Operational Research*, vol. 130, no. 3, pp. 449–467, 2001.

[61] F. B. Hanson, *Applied Stochastic Process and Control for Jump-Diffusions*. Philadelphia: Society for Industrial and Applied Mathematics, 2007.

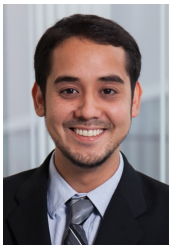
[62] W. Rudin, *Principles of Mathematical Analysis*. McGraw-Hill Education, 3rd ed., 1976.

[63] M. W. Hirsch and H. Smith, *Monotone Dynamical Systems*. 2006.

[64] J. S. Rosenthal, *A First Look at Rigorous Probability Theory*. Hackensack, NJ: World Scientific Publishing Co., 2nd ed., 2013.



Nicholas J. Watkins is currently a Ph.D. student in the Electrical and Systems Engineering program at the University of Pennsylvania. He graduated summa cum laude with a B.S. in Electrical Engineering from Wilkes University in 2013, as the top student in his class. His current research interests are focused on the control and optimization of stochastic systems, with projects focusing on mitigating spread in epidemic processes and managing resources efficiently in energy harvesting devices.



Cameron Nowzari received the Ph.D. in Engineering Sciences from the University of California, San Diego in September 2013. He is currently working as a Postdoctoral Research Associate at the University of Pennsylvania. He was a finalist for the Best Student Paper Award at the 2011 American Control Conference and received the 2012 O. Hugo Schuck Best Paper Award from the American Automatic Control Council. His current research interests include dynamical systems and control, sensor networks, distributed coordination algorithms, robotics, applied computational geometry, event- and self-triggered control, Markov processes, network science, and spreading processes on networks.



George J. Pappas received the Ph.D. in Electrical Engineering and Computer Sciences from the University of California, Berkeley, in December 1998. He is currently the Joseph Moore Professor and Chair of Electrical and Systems Engineering at the University of Pennsylvania. He also holds secondary appointments in Computer and Information Sciences, and Mechanical Engineering and Applied Mechanics. He is a member of the GRASP Lab and the PRECISE Cdenter. He currently serves as the Deputy Dean for Research in the School of Engineering and Applied Science. His current research interests include hybrid systems and control, embedded control systems, cyberphysical systems, hierarchical and distributed control systems, networked control systems, with applications to robotics, unmanned aerial vehicles, biomolecular networks, and green buildings.

0018-9286 (c) 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.