

The Impact of Host-Based Early Warning on Disease Outbreaks

Mark A. Hernandez, Lauren E. Milechin, Shakti K. Davis, Richard A. DeLaura,

Kajal T. Claypool, and Albert J. Swiston

Lincoln Laboratory researchers investigated how early warning of exposure to pathogens could shape health care responses to disease outbreaks. Basing their analysis on the capability of an innovative algorithm that enables the detection of pathogen exposure in individuals before symptoms of disease occur, the research team evaluated the effects of such early warning on various strategies for mitigating a widespread outbreak. While not specifically focused on COVID-19, this effort lays the foundation for understanding the impact of early warning technology in combination with other nonpharmaceutical interventions.



The detection of communicable pathogens responsible for major outbreaks often relies on health care professionals' recognition of symptoms manifesting in infectious individuals. Early warning of such communicable diseases before the onset of symptoms could improve both patient care and public health responses. MIT Lincoln Laboratory's PRESAGED (Presymptomatic Agent Exposure Detection) algorithm uses host-based physiological signals to detect an individual's exposure to pathogens, such as viruses and bacteria, before overt symptoms emerge and infectiousness is peak.

Researchers in the PRESAGED program have used data from non-human primate studies to demonstrate that the algorithm can provide two to three days of early warning before the onset of incipient symptoms (e.g., fever), independent of the particular pathogen, exposure route, pathogen dose, or animal species [1]. These results are consistent with recent findings of Speranza et al. that show presymptomatic upregulation in biomarkers potentially linked to pathogen exposure around the same time in non-human primates exposed to Ebola [2].

The PRESAGED algorithm relies on data that could be collected noninvasively, such as heart electrical activity (electrocardiography, or ECG), to predict the probability that an individual has been exposed to a pathogen. This system addresses fundamental limitations of existing highly sensitive and specific bioassays: specifically, that these tests require the collection of blood or other biofluids and that samples are only collected and analyzed when there is a high suspicion of exposure. Furthermore, most assays focus on detecting proteomic or genetic signatures

of a pathogen itself, at which time an individual will likely already have symptoms and thus be infectious [3].

PRESAGED offers potential applications in health care responses to emerging pathogens on two fronts. The first would be improved individual patient care; medical countermeasures are nearly always more effective when deployed before overt indications of infection [4–6]. The second would be faster and more confident implementation of public health measures to mitigate the spread of disease outbreaks. Lincoln Laboratory researchers have been focusing on the latter, investigating how host-based early warning capabilities could lead to effective pathogen surveillance systems and public health interventions. Although much work has focused on predicting disease outbreaks based on environmental and vector-related data, such as weather patterns, geography, and vegetation conditions [7, 8], less work has explored using the conditions of the host itself for predictive purposes.

This research is of particular interest to the Department of Defense because U.S. servicemembers deployed worldwide can potentially come in contact with a huge variety of communicable diseases, ranging from rarely encountered Ebola to annual influenza. Previously unknown emerging pathogens, such as Chikungunya and Zika, have the potential to disproportionately affect military personnel who are often deployed to regions where such pathogens are prevalent.

Numerous papers have applied epidemiological models to characterize the disease transmission dynamics and the effectiveness of public health interventions of past outbreaks [9–13]. A few papers have incorporated a hypothetical early diagnostic capability (based on advancements in bioassay tests) in their models and evaluated its effectiveness in controlling the spread of disease outbreaks [14]. However, there has yet to be an effort that assesses the potential impact of a host-based early warning system.

We have developed a series of epidemiological models that quantify the potential impact that a host-based early warning capability would have in mitigating pathogen transmission during an outbreak. Our epidemiological models reflect a variety of traditional public health policies related to nonpharmaceutical interventions, such as quarantine and patient isolation, and novel policies enabled by host-based early warning capabilities. For each policy-dependent model, we simulated outbreak scenarios

and calculated the size of the outbreak (total number of infections) and the operational burden (total number of lost duty days resulting from quarantine or isolation). These metrics were used to understand the trade space for the different policies. Our simulations demonstrate the utility of host-based early warning systems in controlling an outbreak under various outbreak conditions, further motivating discussion about the potential benefits and limitations of population-wide implementation in the U.S. armed forces or the broader civilian population.

Models for Disease Outbreak Simulation

The SEIR Model

One of the most common epidemiological models for simulating disease outbreak scenarios is the deterministic SEIR model. This approach splits a given population into separate compartments defined by their relationship to a disease outbreak [15]:

1. Susceptible: healthy individuals who can be exposed to the pathogen
2. Exposed: individuals who are in the incubation phase; they have been exposed to the pathogen but are not yet showing symptoms and are not infectious
3. Infectious: individuals who are infectious to the susceptible population and will eventually display overt symptoms
4. Recovered: those who have recovered from illness and acquired immunity to further infection

These compartments are then linked with a system of ordinary differential equations (ODEs) to characterize how individuals transition into and out of each compartment over time. A variety of scenarios can be simulated by changing the rate parameters of the ODEs linking the population compartments. Furthermore, this approach allows quantitative projections of how many people are exposed to a pathogen and become sick under different outbreak conditions.

While the SEIR model is mathematically rigorous and often has good predictive utility during an outbreak, the SEIR model requires several assumptions. First, the model assumes there is a fixed population N , with no births or deaths other than those resulting from the infectious disease. The next assumption is that the population is homogeneously mixing, meaning that transmission between any two individuals is equally likely. The model also assumes that exposed individuals become infectious

after a fixed incubation period, thus not accounting for individual variability of disease progression for the young, elderly, or immunocompromised. Finally, it is assumed that all recovered individuals are immune to further infection and thus do not re-enter the susceptible class. While any of these assumptions may fail to hold in particular contexts, abiding by them allows for greater mathematical tractability and offers similar relative output trends.

Baseline SEIR Model

The SEIR model consists of a system of four ODEs that describe the rate of change of individuals in each compartment over time (Figure 1). The differential equations for Susceptible, S_t , Exposed, E_t , Infectious, I_t , and Recovered, R_t , are as follows, where dependence on time, t , is omitted in the notation below for simplicity:

$$\begin{aligned} \frac{dS}{dt} &= -\frac{\beta SI}{N} \\ \frac{dE}{dt} &= \frac{\beta SI}{N} - \sigma E \\ \frac{dI}{dt} &= \sigma E - \gamma I \\ \frac{dR}{dt} &= (1 - f)\gamma I \end{aligned}$$

Table 1 describes all of the parameters we use for our models. At the start of the model, a subset of the population is exposed to a pathogen (E_0), and the remainder are susceptible to infection. Susceptible individuals enter the exposed compartment at a rate of $\beta I/N$, which is known as the normalized transmission rate. The parameter β is the contact rate, which accounts for how often susceptible-infectious contacts result in a susceptible individual becoming exposed to the pathogen. Exposed individuals become infectious at a rate of σ , which is the inverse of the incubation period. Infectious individuals stay infectious at a rate of γ (the inverse of the infectious period) until they recover or die. We define the recovery rate, ρ , as $(1 - f)\gamma$ where the case fatality rate f is the proportion of infected individuals who die from the disease. Note that we do not specify a mortality compartment in this analysis, though for some pathogens spread through the handling of remains (such as Ebola during the 2014 West Africa outbreak), this compartment would be a critical addition.

Another important value given by this model is the basic reproduction number \mathcal{R}_0 , which is often included as a basic property of any given pathogen [15]. \mathcal{R}_0 represents the average number of additional infections caused by each infectious individual, assuming there are no control interventions. In a fixed population represented in the SEIR model, the reproduction number can be calculated from parameters in the model as $\mathcal{R}_0 = \beta/\gamma$.

The solutions of the SEIR model are functions of the number of individuals in each compartment with respect to time, i.e., S_t , E_t , I_t , and R_t . We solve this numerically by using an ODE solver in MATLAB®, and with these results, we may now answer some of the initial questions posed. For example, if we want to know the number of people who are not showing symptoms (visibly healthy individuals), we could plot $S + E + R$.

An example output of the full SEIR model is shown in Figure 2a. Here, we begin with a population of $N = 1000$ people (roughly equivalent to a large battalion), and consider a scenario in which 300 individuals are exposed to some pathogen at time, $t = 0$ ($S_{t=0} = 700, E_{t=0} = 300, I_{t=0} = R_{t=0} = 0$). If we assume a particular infectious pathogen has a contact rate, $\beta = 0.75$, with an incubation period of two days, $\sigma = 0.5$, and an infectious period of five days, $\gamma = 0.2$ (similar properties to a highly contagious flu virus), the solution to the SEIR model shows that nearly the entire population will contract the disease over the course of a month. Figure 2b focuses on the number of cumulative infections, a metric for assessing the overall size of the outbreak over time. This baseline condition, in which no public health policy is in effect, represents the most severe outcome from an infectious disease outbreak.

Policy-Dependent SEIR Models

In the previous section, we considered the standard SEIR model, which describes disease spread *without* any public health interventions. To quantitatively evaluate an outbreak scenario with additional measures, the parameters and compartments of the SEIR model may be modified to reflect policy choices or new early warning technologies. Two policies currently considered as standards in handling possible outbreaks rely on self-monitoring to implement voluntary isolation or quarantining all individuals that may have been exposed

TABLE 1. Parameter Definitions and Values for Baseline and Policy-Dependent Models

| PARAMETER | DEFINITION | VALUE FOR THIS ANALYSIS |
|---------------|--------------------------------|-------------------------------------|
| N | Total population | 1000 |
| E_0 | Initially exposed population | 300 |
| β | Contact rate | 0.75 |
| σ | Incubation rate | 0.5 (2-day incubation period) |
| γ | Infectious rate | 0.2 (5-day infectious period) |
| ρ | Recovery rate | 0.18 |
| f | Case fatality rate | 0.1 |
| θ | Self-reporting probability | 1 (all ill individuals self-report) |
| α | Self-reporting rate | 0.5 (2-day self-report) delay |
| T_Q | Time of quarantine | 0 |
| ε | Release rate | 0.2 (5-day quarantine period) |
| κ | False-alarm rate | 1/365 (1 false alarm per year) |
| λ | Early detection rate | 1 (1-day early detection time) |
| P_D | Probability of early detection | 0.8 |

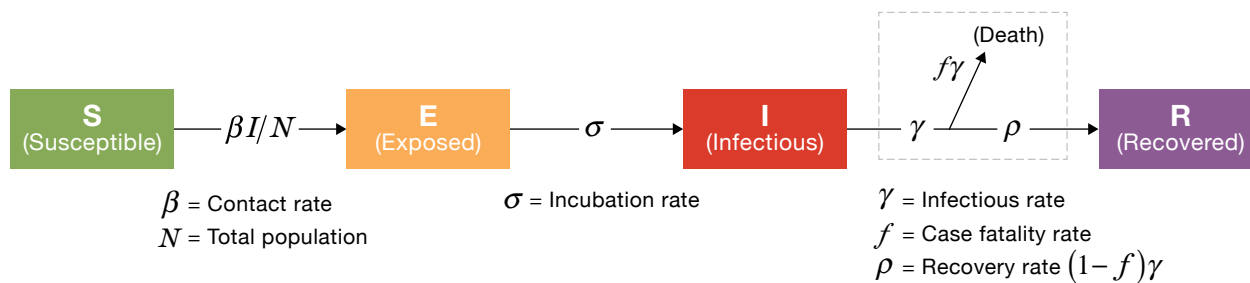


FIGURE 1. In this schematic diagram of the SEIR model, each box represents a compartment of the population, and arrows show the progression of individuals through those compartments. Expressions above arrows are rate coefficients showing progression through compartments. Note that the dashed boxed portion, indicating individuals leaving the population because of death, will not be included in future diagrams for simplicity.

to a pathogen. Isolation applies to individuals who are already ill, whereas quarantine applies to individuals who may have been exposed but have not shown symptoms. The first policy, self-monitoring and reporting, which assumes individuals self-report when they develop symptoms, allows for additional infections during the time delay between the onset of symptoms and isolation;

however, it does not require the logistical and financial cost of quarantine. The second policy, quarantining the entire population, is effective in reducing additional transmissions but requires prohibitively onerous costs for large populations. In an effort to explore an alternative to these two policies, we modified our SEIR model to simulate a third policy, quarantine-on-alert, in which




FIGURE 2. A numeric solution for the baseline SEIR model shows the population of each compartment versus time (a), which indicates that for an initial exposure of 300 individuals, nearly the entire population will eventually fall ill; the cumulative number of infections (Cumulative I) versus time is shown in (b).

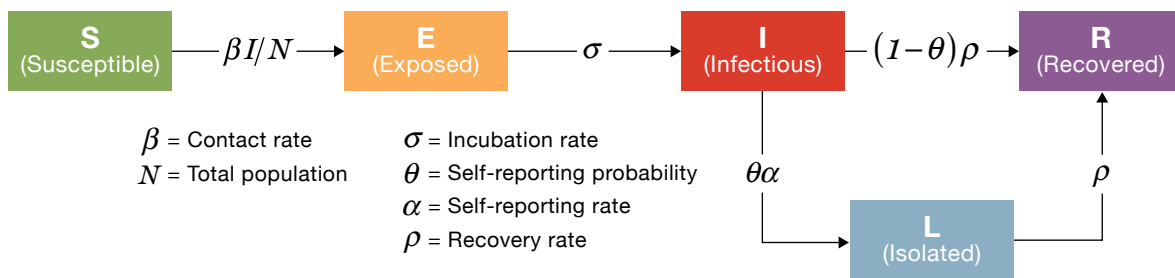


FIGURE 3. The model diagram for a self-monitoring and self-reporting policy includes an additional compartment (L) for individuals who are isolated after developing symptoms. Note that fatalities are left out for simplicity.

individuals are only quarantined when prompted by a host-based early warning system. We hypothesized that the early notification of incipient illness will allow for reductions in pathogen transmission while minimizing the number of quarantined individuals. We then assessed each of these three new policy-dependent SEIR models.

Self-Monitoring and Reporting: Isolation after Symptoms

In the course of an infectious disease outbreak, most individuals will self-monitor for symptoms of the pathogen. If individuals start to develop symptoms, they can self-report to a medical facility and may be

immediately isolated until they recover. A new compartment, isoLated (L), was added to the model to reflect this symptomatic and infectious population, which has limited contact with the rest of the population and therefore reduced transmission potential (Figure 3). The self-reporting probability θ is defined as the proportion of symptomatic individuals who are compliant in reporting their symptoms and enter the L compartment, while the self-reporting rate α is derived from the delay between developing symptoms and self-reporting (i.e., the inverse of the average time delay). For our model, we assumed that isolation is 100 percent effective in preventing transmission, and individuals in the L compartment

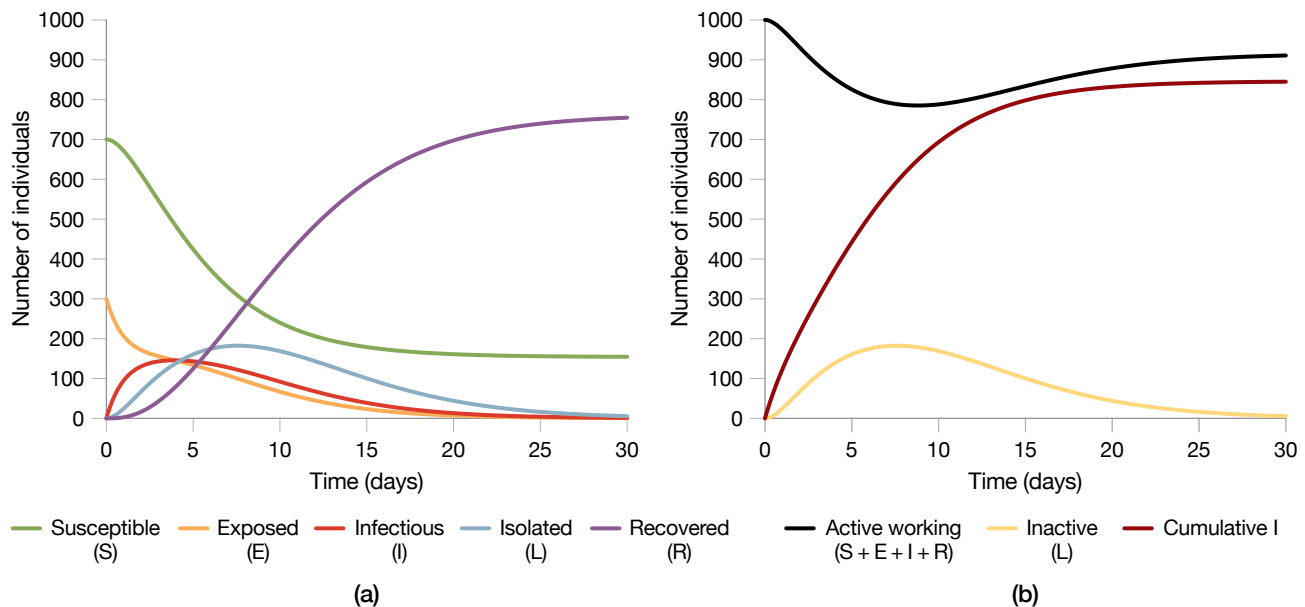


FIGURE 4. In this numeric solution for our SEIR model with a self-monitoring and isolation policy enacted, the population of each compartment versus time is shown in (a), and the active working ($S + E + I + R$) and inactive (L) populations versus time with the cumulative infections versus time are shown in (b). The self-monitoring policy has effectively reduced the total number of disease cases and blunted the outbreak's peak severity (number of cases in I and L peak at approximately 7 days).

are unable to infect the susceptible population. While a self-monitoring policy reduces transmission rates by isolating the self-reported sick, both the latency α and an imperfect self-reporting probability ($\theta < 1$) will still lead to opportunities for transmission.

Outputs for the SEIR model with self-reporting and isolation are shown in Figure 4. This policy addition leads to both fewer total individuals contracting the disease and fewer individuals acutely symptomatic at the height of the outbreak. The new isolated compartment L , however, has other costs associated with lost duty days, mandated isolation, medical facilities, and treatment. The choices of self-reporting probability and rate are critical in this scenario, and, in reality, may have such broad distributions that relying on this policy in an acute outbreak may do little to prevent additional infections.

Quarantining the Entire Population

Another possible, though draconian, public health response to an outbreak is to quarantine currently healthy individuals with some likelihood of pathogen exposure. In outbreaks with a large population subset having some exposure likelihood, a quarantine-all policy would separate all individuals in the population from

contact with one another. Our model assumes that quarantined individuals, represented by the Q compartment, are monitored and immediately isolated once they develop symptoms (Figure 5). If quarantined individuals do not become symptomatic after the maximum incubation time of the suspected pathogen (often several weeks), they are released and re-enter the S compartment at a rate of ϵ (the inverse of the maximum incubation time). As shown in Figure 6, the quarantine-all policy eliminates opportunities for further pathogen transmission; however, it also results in the quarantine of healthy individuals who have not been exposed, contributing to immense productivity losses, extreme logistic burdens associated with providing accommodations for an entire population, and acute civil rights concerns. Furthermore, the assumption that transmission is zero within the Q compartment may not be realistic because transmission for some pathogen may occur before overt symptoms and isolation. This quarantine-all policy could result in the illness of people who, if not for the quarantine, would never have been exposed to the pathogen. Blanket quarantine policies are extreme examples of public health interventions and represent a clear scenario that new technologies may improve.

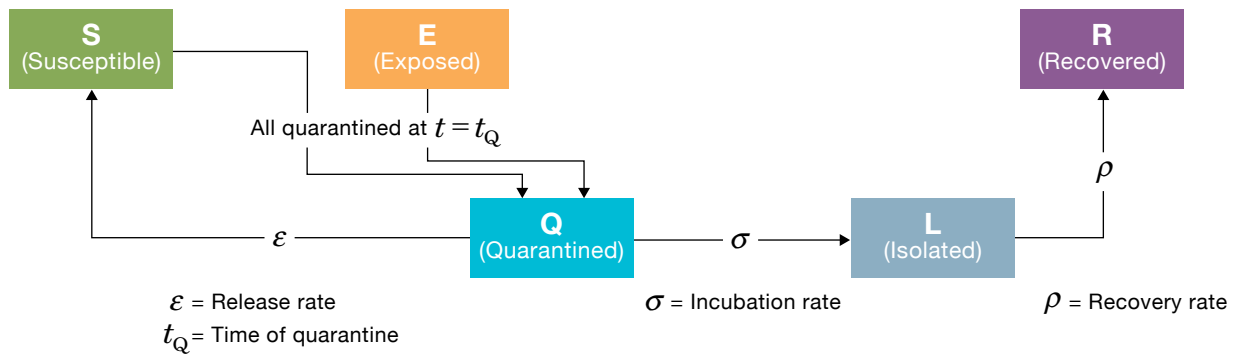


FIGURE 5. The model diagram for a quarantine-all policy moves all susceptible and exposed individuals to the quarantined compartment (Q) at the time $t_Q = 0$ days. Quarantined individuals who do not develop symptoms are returned to the susceptible compartment after a maximum incubation time.




FIGURE 6. In the numeric solution for our SEIR model with a quarantine-all policy enacted, the population of each compartment versus time, which notably does not include an exposed compartment, is shown in (a). In (a) are individuals suspected to be exposed and quarantined initially, until they either fall ill (and are isolated) or are released after the quarantine duration (21 days). The active working (S + R) and inactive (Q + L) populations versus time with the cumulative infections versus time are shown in (b). The quarantine-all policy has very effectively reduced the total number of disease cases but has vastly increased the burdens of quarantine and isolation.

Quarantine on Alert

The third model we explored simulates a novel policy that could potentially be enabled by future host-based early detection capabilities. Here, a host-based early warning system could prompt a quarantine only on alert when the system detects presymptomatic signs of exposure. We integrated the early warning performance metrics from our PRESAGED algorithm—population-wide probability of early detection, daily false-alarm rate, and early

detection time—into a base SEIR model. As shown in Figure 7, a host-based early warning system would trigger exposed individuals to transition into quarantine at a rate characterized by the product of the parameters λ and P_D . The parameter λ is an early detection rate defined as the inverse of the system’s average early detection time for a given pathogen. The parameter P_D is the population-wide probability of early detection, i.e., the fraction of the exposed population that will present early detections of

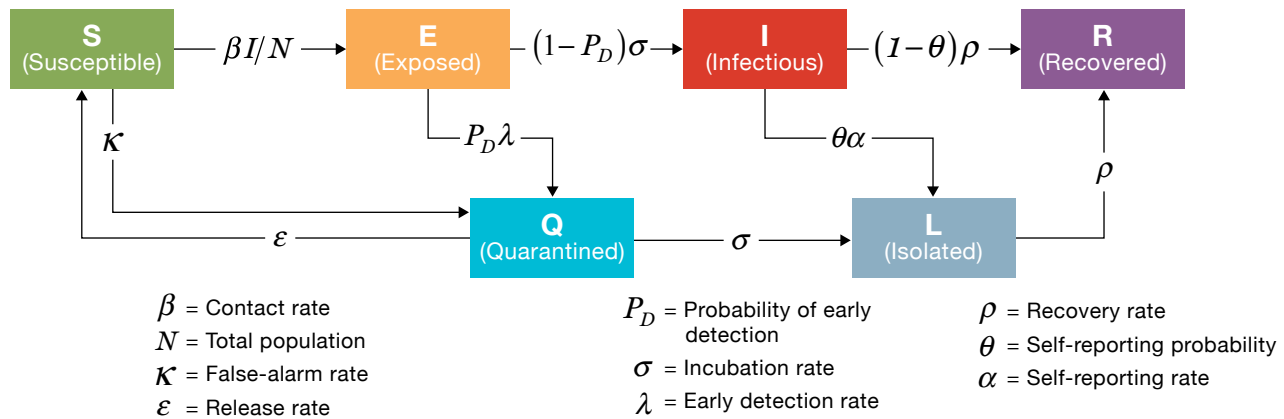


FIGURE 7. This hybrid model diagram illustrates the final quarantine-on-alert policy. The performance of the PRESAGED-like early warning system is now explicitly added, including the probabilities of early detection and false alarm, as well as an expected early detection time before symptom onset and infectiousness. This hybrid between quarantine-all and self-monitoring with isolation seeks to leverage the epidemic-limiting behavior of both policies while reducing the cost and burdens of each.

pathogen exposure prior to becoming infectious. Missed detection cases, when the system fails to produce an early warning, occur with a probability of $1 - P_D$ across the exposed population. For this policy, false alarms of the host-based early warning system would cause healthy individuals to be incorrectly quarantined. Our model takes this notion into account by adding a transition rate from the susceptible to the quarantined compartment, characterized by a daily false-alarm rate κ , the probability of an individual’s early warning system presenting a false alarm on a given day.

While the self-monitoring and quarantine-all policies present clear trade-offs between cumulative number of infections and number of people isolated or quarantined, as shown in Figure 8, this quarantine-on-alert policy offers a hybrid approach that could optimize for both. Quarantine-on-alert vastly reduces quarantine costs by selectively identifying likely exposed individuals through physiologically based predictions. Furthermore, once quarantined, these high-risk individuals can be immediately isolated, thus limiting the opportunities of additional infections and enabling more focused medical care.

SEIR Model Reproduction Numbers by Policy Choice

Modifications made to the baseline SEIR model result in changes to the reproduction number (\mathcal{R}) associated with each policy-dependent SEIR model [16]. Our example scenario assumes that every symptomatic person will

self-report and thus the self-reporting probability, θ , is 1. Table 2 summarizes equations to calculate \mathcal{R} values for each policy, the \mathcal{R} values given the parameter values in Table 1, and the cumulative cases of infection for each policy choice after 50 days. The policies with the lower \mathcal{R} values (quarantine-all and quarantine-on-alert) are associated with minimal pathogen spread, while the policies with higher \mathcal{R} values (self-monitoring and baseline) yield a greater number of infection cases. The number of cumulative infections as the ODE reaches equilibrium illustrates how the reproduction number would affect the final size of the outbreak. Note that the number of cumulative infections for baseline saturates around the population size ($N = 1000$) when the number of individuals in the susceptible compartment is exhausted.

Policy Trade-Space Analysis

The policy-dependent SEIR models provide a foundation for a trade-space analysis of different quarantine, isolation, and treatment (QIT) policies as a function of disease transmission characteristics, exposure scenarios, and performance of early warning systems. The analysis captures both the potential benefits (reduction in quarantine costs, more focused use of medical resources) and risks (increase in infections) for each policy.

To demonstrate the potential utility of the early warning-enabled quarantine-on-alert model in a more comprehensive QIT policy analysis, we compare outcomes over a range of disease transmission rates,

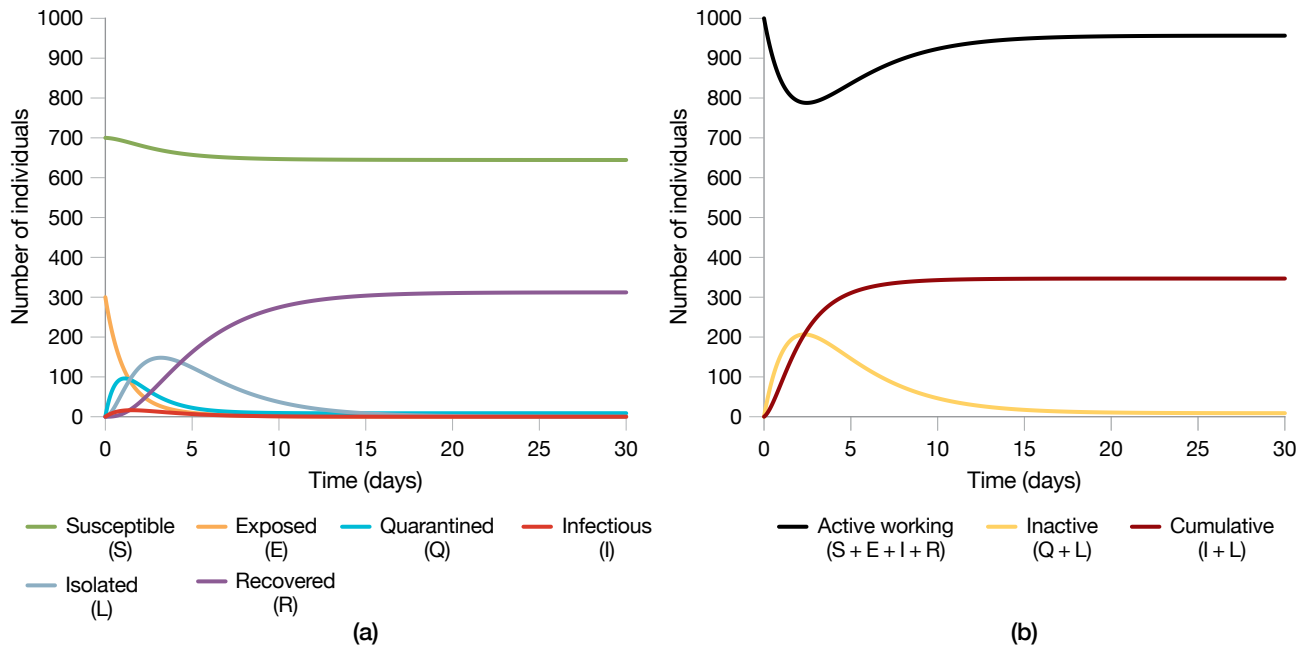


FIGURE 8. This numeric solution for our SEIR model with a quarantine-on-alert policy enacted shows the population of each compartment versus time (a); the plot shows how such an early warning technology could limit the outbreak size. As the initially exposed individuals receive alerts, they enter the quarantine-on-alert compartment and are unable to infect the susceptible compartment. The active working ($S + E + I + R$) and inactive ($Q + L$) populations versus time are shown with the cumulative infections versus time in (b). The quarantine-on-alert policy has effectively reduced the additional disease cases and has largely limited those in quarantine to those individuals who were initially exposed.

initial exposure conditions, and early warning performance parameters. We vary these parameters in simulations of a mass exposure to an infectious pathogen occurring in a population of 1,000 people. All scenarios were simulated over 50 days. Because we fix the exposure time for all cases to $t = 0$, rather than using the quarantine release rate ϵ , all individuals in quarantine are returned to the susceptible compartment at $t = 21$ days (the maximum incubation period).

Policy outcomes are defined by two metrics that are derived from the policy-dependent SEIR model outputs as shown in Figure 9. The first, lost duty days, is expressed as a percentage of total number of days of work productivity that are lost because of quarantine and isolation in the 50-day simulation; this percentage is proportional to the integral of the curves in Figure 9a. The second metric, cumulative infections, is the percentage of the population that has been infected by the end of the simulation (Figure 9b). These two metrics characterize a trade space for evaluating QIT policies under different circumstances, because any measure that reduces the number of people

in quarantine and isolation may be expected to increase the likelihood of infection in the population. While these metrics are helpful for evaluating the impact of various policy choices, they are by no means comprehensive and particularly do not consider the financial or other costs associated with quarantine.

Five QIT policy sets were evaluated: quarantine-all, isolate upon self-reporting, and quarantine-on-alert with three different levels of early detection performance—high sensitivity ($P_D = 0.8$, $\kappa = 0.10$), high specificity ($P_D = 0.4$, $\kappa = 1/365$), and a near-ideal early-detection system ($P_D = 0.8$, $\kappa = 1/365$). These five QIT policy options were then tested against four different outbreak scenarios with high or low disease transmission (where $\beta = 0.6$ or 0.3 , respectively) and high or low initial population exposures (where $E_0 = 600$ or 50 , respectively). Figure 10 shows the QIT policy trade-offs for these 20 independent outbreak and policy combinations.

For all scenarios, a quarantine-on-alert policy coupled with the near-ideal early warning capability reduces quarantine needs with only a small increase in

TABLE 2. Reproduction Number (\mathcal{R}) Equations and Nominal \mathcal{R} Values for Our Example Scenario

| POLICY | \mathcal{R} EQUATION | \mathcal{R} VALUE | CUMULATIVE INFECTIONS AFTER 50 DAYS |
|---------------------|---|---------------------|--|
| Baseline | $\mathcal{R}_0 = \frac{\beta}{\gamma}$ | 3.75 | 986 |
| Self-monitoring | $\mathcal{R} = \frac{\beta}{\alpha}$ | 1.5 | 846 |
| Quarantine-all | $\mathcal{R} = 0$ | 0 | 300 |
| Quarantine-on-alert | $\mathcal{R} = \frac{\beta}{\alpha} \frac{(1 - P_D)\sigma}{(1 - P_D)\sigma + P_D\lambda}$ | 0.167 | 347 |

In a population of $N = 1000$, the number of initially exposed individuals is $E_0 = 300$. The nominal \mathcal{R} Value and Cumulative Infections were calculated using parameter values from Table 1.

the number of additional infections. The cost of focusing on a high-specificity early detection system (i.e., a reduction in false alarms and thus quarantine costs) is an increase in additional infections relative to the near-ideal system. Conversely, a high-sensitivity system increases the percentage of the population in quarantine compared to both the ideal and high-specificity early detection system while also reducing the number of additional infections to nearly the number seen by quarantining the entire population a priori.

For scenarios with low initial exposure, the impact of early warning-enabled QIT policies varies dramatically with transmissibility. For low-transmissibility and small initial exposures (Figure 11a), an early warning capability demonstrates the least utility because the outbreak seems to be effectively contained with just self-monitoring. However, for a pathogen with high transmissibility (Figure 11b), the cumulative infections are reduced significantly with any of the notional early warning systems. A high-specificity system offers a particularly promising result, showing comparable infection reduction to the near-ideal system while also minimizing the number of days lost to quarantine.

In the scenarios of high initial exposure (Figure 11c and Figure 11d), the costs of imperfect early detection are more pronounced compared to those for scenarios of low initial exposure. Because a high initial exposure greatly increases the likelihood of infectious individuals

transmitting the pathogen to those susceptible, even a short delay in isolating infectious individuals will lead to more infections. The reduction of cumulative infections for all notional early warning systems, relative to the self-monitoring policy, is less prominent than in scenarios of low initial exposure because the initial exposed population is close to saturation. In effect, for these extreme cases of mass exposures, the more aggressive policies (i.e., a high-sensitivity early detection system or quarantining everyone) may be more effective.

Discussion and Future Work

The results for the early warning-enabled and policy-dependent SEIR models allowed for a quantitative analysis of the QIT trade space and provided guidance on priorities for the future development of early warning technology. However, this illustration of QIT risk analysis captures only a subset of the factors that must be considered in the formulation of a rational, effective QIT policy. The context of the scenario will ultimately inform where the ideal operating point would be. For example, the number of total infections may be interpreted differently depending upon the virulence of the disease. The tolerance for new infections may be low if the consequences are high, such as if the infection is almost always fatal or is associated with severe symptoms and long-term complications. Additionally, the availability of diagnostic tests may further refine the use of early warning capabilities;

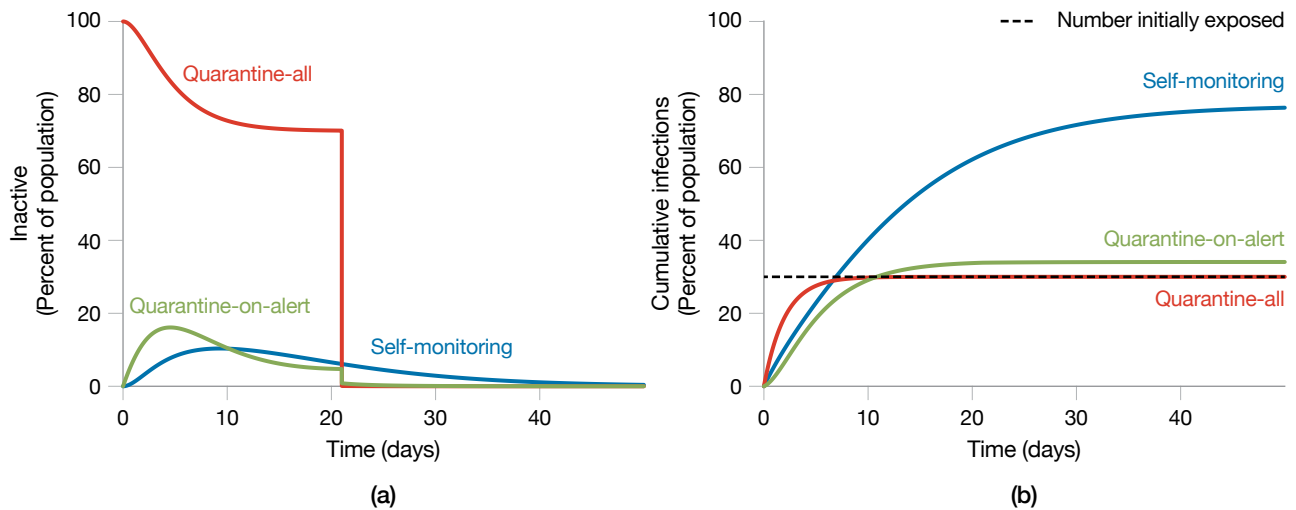


FIGURE 9. The model outputs compare consequences of self-monitoring, quarantine-all, and quarantine-on-alert outbreak response policies. In (a), the plot shows the number of inactive individuals who are unable to work because of quarantine and isolation ($Q + L$ for quarantine-all and quarantine-on-alert, only L for self-monitoring because it does not utilize quarantine) versus time. The sharp drop-off for quarantine-all represents the release of all quarantined individuals after 21 days. In (b), the plot shows the cumulative infections versus time. Any infections above the dotted line (number of individuals initially exposed at $t = 0$) indicate additional transmission of the pathogen.

a more sensitive early detection capability could be combined with a cued use of diagnostic testing to form a much more targeted approach, reducing costs and the likelihood of a false detection. While the results we presented are a first attempt at understanding the potential utility of a host-based early warning system during an outbreak, a comprehensive risk assessment of a QIT policy must consider a range of factors:

- Disease characteristics: prevalence, transmissibility, incubation period, and severity
- Response options: reliable diagnostic tests, vaccines, or treatments
- Resources: cost and availability of QIT measures
- Early warning capability: performance characteristics of early warning systems
- Public health infrastructure: ability to implement an effective public health campaign

An element that we do not include in this model is patient care measures, which are nearly always more effective when deployed earlier than overt indications of infection (e.g., fever). For instance, antiviral drugs (such as zanamivir and oseltamivir/Tamiflu) are most effective in the first ~48 hours of symptoms [8, 17, 18]; PRESAGED-enabled early warning would allow much faster prescription, use, and potentially more profound

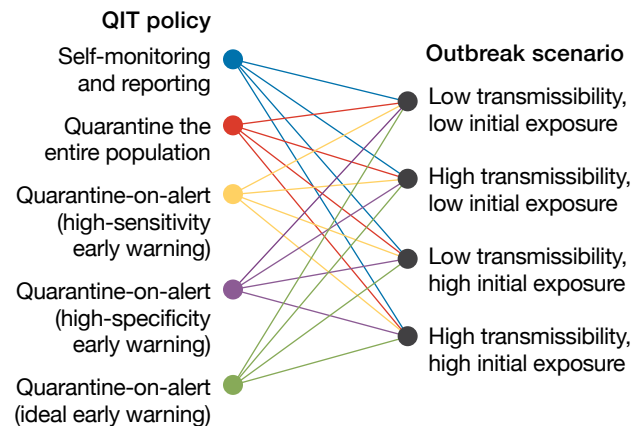


FIGURE 10. Simulations are run on each combination of QIT policy and outbreak scenario, resulting in 20 independent conditions for the trade-space simulations.

therapeutic impact for current dosage recommendations. Triggering the use of diagnostics early would allow clinicians to target drug-based interventions, such as using the precise antibiotic for bacterial infections rather than relying on broad-spectrum versions that contribute to the evolution of drug-resistant bacteria (“superbugs”). Finally, simple supportive care for mild infections, e.g., having rehydrating solutions or over-the-counter

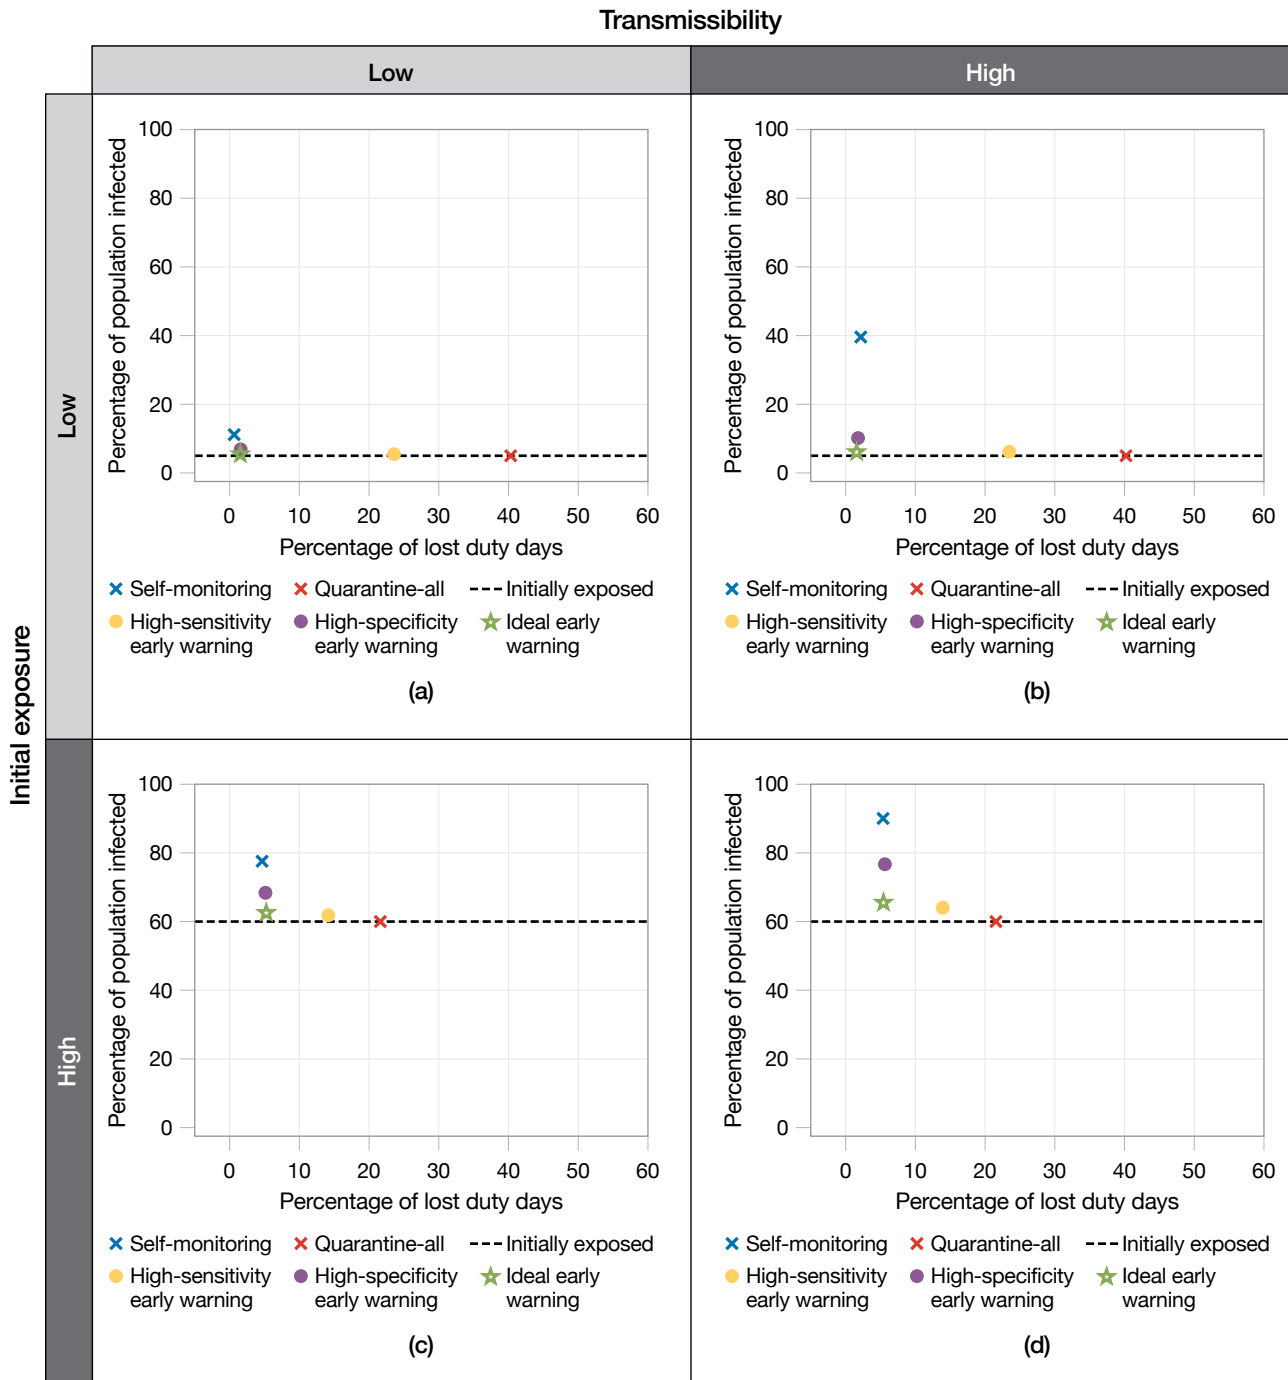


FIGURE 11. The trade space for the percentage of population infected versus percentage of lost duty days compares the augmented policies: self-monitoring, quarantine-all, and quarantine-on-alert with three levels of early warning performance (high sensitivity, high specificity, and ideal early warning). Four situations were considered: low-transmissibility pathogen with a low initial exposure (a), high-transmissibility pathogen with a low initial exposure (b), low-transmissibility pathogen with a high initial exposure (c), and high-transmissibility pathogen with a high initial exposure (d).

symptom-relieving medications available, could ease the burden on health care workers in hospitals, nursing homes, or college campuses. All of these exciting possibilities could be enabled or improved with earlier detection.

Future modeling efforts will focus on improvements to this model, especially as more is known about the mechanistic basis for host-based early warning of pathogen exposure. Several additional aspects of QIT policy that will be considered include the following:

1. Consequences of infection, or a cost function of being in the infectious compartment. Adding a more explicit fatality compartment is straightforward. However, more complex models incorporating significantly time-delayed or nonlinear costs are much more accurate for measuring the effectiveness of QIT policies.
2. Early quarantine release. When the maximum incubation period may result in long quarantine periods that lead to unacceptably high quarantine costs, the absence of an exposure detection in individuals who have not been exposed to the pathogen may trigger an early release from quarantine. This early quarantine release approach could significantly reduce costs and civil rights issues associated with quarantine.
3. Modeling of additional QIT responses. Incorporating diagnostic tests into the models could refine the quarantine trigger and release policies. Additionally, future models could include specific treatments upon early warning to modulate infection outcomes. Such options may be particularly important in circumstances when quarantine and isolation resources are limited.
4. Movement among populations. The policy-dependent SEIR model addresses a single, isolated population being homogeneously mixed. However, in many circumstances, QIT policy is driven by a concern for the spread of the pathogen into connected populations. To address these circumstances, the policy-dependent SEIR model may be further extended to incorporate changes in each population compartment that result from the movement of people into (or out of) individual population compartments from (or to) other populations. In this formulation, individual populations are modeled with the extended SEIR mode as nodes in a migration network. Network edges are characterized by migration rates between the connected subpopulations.

5. Stochastic modeling. In a more realistic model, each of the SEIR parameters may be modeled in greater detail as a probability distribution rather than a fixed value. Stochastic modeling would support a more comprehensive risk analysis that includes insights about the likely range of potential outcomes, as well as rare but high consequence “edge” cases.

In conclusion, we have shown the epidemiological value of host-based early warning systems in a variety of pathogen outbreaks. By adjusting the underlying assumptions, both of the outbreak and the system performance metrics of a notional early detection system, we show in which scenarios early detection is most impactful. The results of this work emphasize the value of early detection in modulating public health responses, though future efforts will also include the value to individual patients. Current efforts at Lincoln Laboratory are focusing on the ability to monitor pathogen exposure of annual influenza, the ability to measure pathogen transmissibility non-invasively, and the impact this detection capability will have on patient care, public health responses, and service member readiness. ■

References

1. L. Milechin, S. Davis, T. Patel, M. Hernandez, G. Ciccarelli, S. Schwartz, et al., “Detecting Pathogen Exposure during the Non-symptomatic Incubation Period Using Physiological Data,” *bioRxiv*: 218818-218818, 2017.
2. E. Speranza, S.L. Bixler, L.A. Altamura, C.E. Arnold, W.D. Pratt, C. Taylor-Howell, et al., “A Conserved Transcriptional Response to Intranasal Ebola Virus Exposure in Nonhuman Primates Prior to Onset of Fever,” *Science Translational Medicine*, vol. 10, no. 434, 2018.
3. R. Liu, X. Wang, K. Aihara, and L. Chen, “Early Diagnosis of Complex Diseases by Molecular Biomarkers, Network Biomarkers, and Dynamical Network Biomarkers,” *Medicinal Research Reviews*, vol. 34, no. 3, 2014, pp. 455–478.
4. P.K. Tosh and P. Sampathkumar, “What Clinicians Should Know about the 2014 Ebola Outbreak,” *Mayo Clinic Proceedings*, vol. 89, no. 12, 2014, pp. 1710–1717.
5. D.G. Bausch, C.M. Hadi, S.H. Khan, and J.J. Lertora, “Review of the Literature and Proposed Guidelines for the Use of Oral Ribavirin as Postexposure Prophylaxis for Lassa Fever,” *Clinical Infectious Diseases*, vol. 51, no. 12, 2010, pp. 1435–1441.
6. G. Stiver, “The Treatment of Influenza with Antiviral Drugs,” *Canadian Medical Association Journal*, vol. 168, no. 1, 2003, pp. 49–56.
7. T.E. Ford, R.R. Colwell, J.B. Rose, S.S. Morse, D.J. Rogers, and T.L. Yates, “Using Satellite Images of Environmental

- Changes to Predict Infectious Disease Outbreaks,” *Emerging Infectious Diseases*, vol. 15, no. 9, 2009, pp. 1341–1346.
8. S. Altizer, R.S. Ostfeld, P.T. Johnson, S. Kutz, and C.D. Harvell, “Climate Change and Infectious Diseases: From Evidence to a Predictive Framework,” *Science*, vol. 341, 2013, pp. 514–519.
 9. J. Legrand, R.F. Grais, P.Y. Boelle, A.J. Valleron, and A. Flahault, “Understanding the Dynamics of Ebola Epidemics,” *Epidemiology and Infection*, vol. 135, no. 4, 2007, pp. 610–621.
 10. A.J. Kucharski, S. Funk, R.M. Eggo, H.P. Mallet, W.J. Edmunds, and E.J. Nilles, “Transmission Dynamics of Zika Virus in Island Populations: A Modelling Analysis of the 2013–14 French Polynesia Outbreak,” *PLOS Neglected Tropical Diseases*, 2016, pp. 1–15.
 11. C.E. Mills, J.M. Robins, and M. Lipsitch, “Transmissibility of 1918 Pandemic Influenza,” *Nature*, vol. 432, no. 7019, 2004, pp. 904–906.
 12. G. Chowell and H. Nishiura, “Transmission Dynamics and Control of Ebola Virus Disease (EVD): A Review,” *BMC Medicine*, vol. 12, no. 196, 2014, pp. 1–16.
 13. A. Flahault, E. Vergu, L. Coudeville, and R.F. Grais, “Strategies for Containing a Global Influenza Pandemic,” *Vaccine*, vol. 24, no. 44–46, 2006, pp. 6751–6755.
 14. D. Chowell, C. Castillo-Chavez, S. Krishna, X. Qiu, and K.S. Anderson, “Modelling the Effect of Early Detection of Ebola,” *The Lancet Infectious Diseases*, vol. 15, no. 2, 2015, pp. 148–149.
 15. O. Diekmann and J. Heesterbeek, *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*. Chichester, U.K.: Wiley & Sons, 2000.
 16. P. van den Driessche and J. Watmough, “Reproduction Numbers and Sub-threshold Endemic Equilibria for Compartmental Models of Disease Transmission,” *Mathematical Biosciences*, vol. 180, no. 1–2, 2002, pp. 29–48.
 17. A.S. Monto, D.M. Fleming, D. Henry, R. DeGroot, M. Makela, T. Klein, et al., “Efficacy and Safety of the Neuraminidase Inhibitor Zanamivir in the Treatment of Influenza A and B Virus Infections,” *The Journal of Infectious Diseases*, vol. 180, no. 2, 1999, pp. 254–261.
 18. F.G. Hayden, A.D. Osterhaus, J.J. Treanor, D.M. Fleming, F.Y. Aoki, K.G. Nicholson, et al., “Efficacy and Safety of the Neuraminidase Inhibitor Zanamivir in the Treatment of Influenzavirus Infections,” *New England Journal of Medicine*, vol. 337, no. 13, 1997, pp. 874–880.


About the Authors



Mark A. Hernandez is a member of the technical staff in the Human Health and Performance Systems Group. Since joining Lincoln Laboratory in 2015, he has worked on building data pipelines and predictive models for applications in physiological signal analysis, biomedical imaging, and early detection of infectious disease. His research interests focus on developing data-driven technologies that address issues in public health. He holds a bachelor’s degree in electrical engineering from the University of Illinois at Urbana-Champaign and is currently pursuing a master’s degree in public health at Boston University through the Lincoln Scholars Program.



Lauren E. Milechin is a member of the technical staff in the MIT Department of Earth, Atmospheric, and Planetary Sciences. Her interests and projects involve big data, database technology, machine learning, and applications of these topics to other fields. She also assists researchers on the MIT SuperCloud supercomputing system. Previously, she worked as a technical staff member at the Lincoln Laboratory Supercomputing Center. She holds a bachelor’s degree in mathematical sciences from Worcester Polytechnic Institute, where she explored mathematical modeling of disease and of population dynamics, and a master’s degree in industrial mathematics from the University of Massachusetts Lowell, focusing on computer science applications, such as machine learning and algorithms.



Shakti K. Davis is a member of the senior technical staff in the Human Health and Performance Systems Group. Among her research interests is the development of sensor and signal processing algorithms that are aimed at applications for protecting warfighters against injury and enhancing warfighters’ performance capabilities. Since 2015, she has supported a variety of signal processing research efforts, including characterizing exposure to hazardous military noise through the use of body-worn acoustic sensors, developing algorithms to estimate risk factors of musculoskeletal injury by using foot-mounted biomechanical sensors, and investigating classification methods that use high-resolution noninvasive physiological sensors to provide early warning of virus exposure. She joined Lincoln Laboratory in 2006 and worked on detection and classification algorithms for airborne radar systems for nine years. She holds a bachelor’s degree from New Mexico State University and a doctorate from the University of Wisconsin–Madison in electrical engineering.



Richard A. DeLaura is a technical staff member in the Air Traffic Control Systems Group. Prior to joining Lincoln Laboratory in 2000, he was a research scientist at the University of Massachusetts Dartmouth, developing software and curriculum to introduce advanced mathematical concepts to middle and high school

students. He has authored or coauthored several conference proceedings and Federal Aviation Administration research reports. He holds a bachelor's degree in chemistry and physics from Harvard University.



Kajal T. Claypool is a member of the senior staff in the Human Health and Performance Systems Group. Her current research focuses on translational informatics to understand the associations between environment, omics, and disease. She is particularly interested in using statistics and machine learning to

investigate the pseudo-intervention of physical activity on obesity, aging, and Alzheimer's disease. She is also actively engaged in research and technology development in big data management and has been active in the database research community for more than 20 years. Her interests in this field include information integration, information retrieval, data analytics, and secure architectures. She has authored numerous publications in these and related areas. Prior to joining Lincoln Laboratory in 2007, she worked as a member of the technical staff at Oracle Corporation and as an assistant professor in the Department of Computer Science at the University of Massachusetts Lowell. She holds a post-doctoral master's degree in biomedical informatics from Harvard Medical School and a doctorate in computer science from Worcester Polytechnic Institute.



Albert J. Swiston was a principal investigator in the Human Health and Performance Systems Group at Lincoln Laboratory. His research focused on novel medical devices and physiological signals analysis for military and civilian applications. His work has been published in *Nature Materials*, *Science Translational*

Medicine, *PLoS* publications, and IEEE journals. He served as a National Security Scholar in the U.S. intelligence community, and as a Science and Technology Policy Fellow at the National Academies' Committee on Science, Engineering, Medicine, and Public Policy. He holds bachelor's and master's degrees from Johns Hopkins University and a doctorate in materials science and engineering from MIT. He is currently the Director of Digital Operations and Innovation at Merck Research Labs.