



OPEN

Simulating nationwide coupled disease and fear spread in an agent-based model

Joy Kitson^{1,2}✉, Prescott C. Alexander^{2,3}, Joseph Tuccillo⁴, David J. Butts³, Christa Brelsford³, Abhinav Bhatele¹, Sara Y. Del Valle⁵ & Timothy C. Germann²✉

Human cognitive responses, behavioral responses, and disease dynamics co-evolve over the course of any disease outbreak, and can result in complex feedbacks. We present a dynamic agent-based model that explicitly couples the spread of disease with the spread of fear surrounding the disease, implemented within the EpiCast simulation framework. EpiCast models transmission within a realistic synthetic population, capturing individual-level interactions. In our model, fear propagates through both in-person contact and broadcast media, prompting individuals to adopt protective behaviors that reduce disease spread. In order to better understand these coupled dynamics, we create and compare a range of compartmental models to ensure that introducing additional disease states does not prevent the emergence of multiple waves in these simpler models. Additionally, we compare a range of behavioral scenarios within EpiCast, varying the level and intensity of fear and behavior change. Our results show that the addition of asymptomatic, exposed, and pre-symptomatic disease states can impact both the rate at which an outbreak progresses and its overall trajectory in compartmental models. In EpiCast, the combination of non-local fear spread via broadcasters and strong behavioral responses by fearful individuals generally leads to multiple epidemic waves, an outcome that occurs only within a narrow parameter range when fear spreads purely through local contact. Accounting for the coupled spread of fear and disease is critical for understanding disease dynamics and designing timely, targeted responses to emerging infectious threats.

Keywords Agent-based model, Coupled contagions, Infectious diseases, Disease propagation

Looking back on the early stages of the COVID-19 pandemic reveals a dizzying array of responses across all levels of society, across the entire world. The response to the COVID-19 pandemic may have been the largest coordinated behavior change in the history of humanity. From state-mandated lockdowns and school closures to individual masking, testing, and vaccination, these responses shaped both personal experiences of the pandemic and observed outcomes at a societal level^{1–4}.

Unfortunately, accurately anticipating individual behavior has proved challenging for the similarly diverse array of models developed to simulate the disease's spread⁵. The most common method of accounting for behavioral changes has been to use exogenous data sources, such as cell phone mobility data or historical public health policies, to impose top-down behaviors⁶. However, exogenous approaches struggle to capture the two-way feedback between behavioral changes and disease spread, since they mostly treat behavior as an input parameter, able to be scheduled in advance⁷. Truly dynamic behaviors require incorporating a behavioral model into the heart of a simulation, able to respond flexibly to changes in the global or local simulation state.

However, most such endogenous behavioral models operate within high-level compartmental models, where a series of ordinary differential equations (ODEs) govern the changing proportion of the population in different disease states^{8,9}. Such models can only answer population-level questions about disease spread because they smooth over the heterogeneity of real-world populations. In contrast, agent-based models (ABMs) capture this heterogeneity by directly modeling the disease state and behavior of each individual in the simulated population. This capability has enabled the use of ABMs for retrospective analysis of COVID-19 spread patterns¹⁰,

¹Department of Computer Science, University of Maryland, College Park, MD 20742, USA. ²Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. ³Analytics, Intelligence and Technology Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. ⁴Geospatial Science and Human Security Division, Oak Ridge National Laboratory, Oak Ridge, TN 37830, USA. ⁵Associate Laboratory Directorate for Global Security, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. ✉email: jkitson@umd.edu; tcg@lanl.gov

evaluation of public health interventions^{11–13}, and the study of differential COVID-19 outcomes across regions and demographic groups¹⁴. However, ABMs rarely include endogenous behaviors⁸, and the few existing ABMs which do only simulate relatively small populations consisting of at most a million people¹⁵.

To address these gaps, we present a coupled contagion model of interdependent disease and fear spread implemented in EpiCast. EpiCast is a production ABM capable of simulating the spread of an infectious disease – such as COVID-19 – on a digital twin of the full US population while accounting for varying demographic and occupational distributions, as well as commute and flight patterns across the country^{11,16}. In our model, fear of a disease spreads alongside the disease itself, influencing individuals to adopt various protective behaviors that either reduce their number of contacts or the likelihood they are infected by any given contact. This fear can spread through either in-person contact – much like the disease – or through broadcast media. To our knowledge, this is the first work to model the coupled dynamics of two contagions in an agent-based simulation of this scale and complexity.

Additionally, we present a series of ODE-based models we developed to ensure that solely introducing new disease states used in EpiCast would not disrupt the dynamics found in existing ODE-based coupled contagion models. EpiCast includes several disease states beyond the standard susceptible, (symptomatic) infectious, recovered (SIR) model, such as asymptomatic and presymptomatic infectious states. As most prior coupled contagion models have been based on the SIR model, we found it useful to first adapt existing ODE-based models to account for these new disease states, settling on a high-level design before translating the resulting system into the agent-based context. We then analyzed how the incremental inclusion of disease states in these ODE models changes their dynamics.

In the process of designing and evaluating our model, we sought to answer the following research questions:

RQ 1: How does the introduction of additional disease compartments change outbreak dynamics?

RQ 2: How do different types of fear-induced behavior impact outbreak dynamics in a large scale ABM?

RQ 3: How do key parameters in our models relate to important outbreak characteristics, such as the total attack rate or the number of epidemic waves?

Related work

The most common mechanism for incorporating changing behavior into epidemic models is using data-driven, *exogenous* behavior models. Exogenous models treat behavior as an independent variable, rather than a simulation output; they do not allow behavior to be tightly coupled with simulation results like the prevalence of the disease at a given time. In contrast, *endogenous* behavior models incorporate a framework for human behavior directly into a simulation. For example, students could stay home from school after their teacher tests positive for a disease. This tight coupling allows endogenous models to react flexibly to novel scenarios. Such models have proven able to produce multiple epidemic waves and at times dramatically improve model predictions^{17,18}, and generally fall into three categories⁸:

1. *Feedback loop models*, where behavior directly depends on disease prevalence. This can be in the form of either agents adopting behavior when the disease prevalence passes a certain threshold, or by scaling adoption probabilities or intensities based on prevalence. One noteworthy example by Rahmandad et al. outperformed the majority of the U.S. Center for Disease Control COVID-19 Forecasting Hub models by scaling the transmission rate in a SEIR compartmental model by a prevalence-dependent scaling factor¹⁷. Kasa et al. provide an overview of such models¹⁹.
2. *Game theoretic models*, where individuals select an option that maximizes a utility function describing their preferred outcomes. These models can typically consider two types of *agents*, individual decision-makers with their own behaviors: 1) central planners, which select policies or incentive structures to maximize a global measure of social good, and 2) self-interested agents seeking to maximize their private utility. These two types of agents may be used alone or together. Huang et al. survey models of this type²⁰.
3. *Coupled contagion models*, where an individual's behaviors depends on their beliefs, or some other internal state. An individual's beliefs are generally represented as discrete compartments, similar to disease states, with transition probabilities determined by a combination of other individuals' beliefs and information regarding the state of the pandemic. Beliefs and prevalence information can both be either global or local, spreading between neighbors in an ABM, and therefore may spread in a manner similar to the disease. Epstein et al. present such a model where two contagions – fear of the disease and the disease itself – spread simultaneously through a population²¹, adding a third contagion – fear of vaccination – in a later version of the model²².

Several models also combine elements from multiple categories above. Poletti et al. developed a behavioral model for H1N1 Influenza where individuals choose between protective and normal behaviors so as to maximize a utility function based on the recent disease prevalence²³. Jovanovic et al. compare models of vaccine uptake using an objective function which incorporates population-level prevalence information and local-level information spread through neighbors in a information diffusion network²⁴.

Agent-based modeling of endogenous behavior

Despite the range of developed behavioral models, the vast majority are built on top of standard compartmental models, most of which add a simple feedback loop to the system of ordinary differential equations (ODEs) for

the model. In particular, there are only a handful of agent based models (ABMs) which use coupled contagion models⁸.

Epstein et al.²¹ develop a relatively simple ABM where 1800 agents move around a 2D grid. Agents can spread fear and disease to agents in adjacent cells, and can either flee, isolate, or retain the same behavior while fearful. Note that while they demonstrate the emergence of waves in the ODE version of their model, they fail to do so for the ABM. Rajabi et al.²⁵ also use a grid-based ABM, but only have agents quarantine when fearful. They demonstrate how multiple waves can emerge in their simulation and simulate the impact of travel restrictions and contact tracing on disease spread in a 5000 agent population. Both ABMs are highly idealized, with agents lacking realistic schedules. Palomo-Briones et al.²⁶ present a more complex but still spatially explicit 400 agent model where agents have a short visit schedule consisting of visits to their home, social activities, and economic activities, and can chose to wear masks, social distance, or wash hands based on a combination of other agents' beliefs and the results of their previous behavior. Curiel et al.²⁷ use a small 5000-agent network-based ABM where a disease spreads alongside anti-vaccine beliefs under different network topologies and vaccination strategies. While these two models introduce more sophisticated behaviors, they deal with very small populations, much like the previous models. Mao¹⁵ presents a detailed simulation of the Buffalo metropolitan area with nearly a million agents, with a schedule of activities drawn from demographic data. Agents can decide to use flu prophylaxis using a threshold-based decision model based on either infection risk or prophylaxis adoption among contacts. While these models range in complexity, only Rajabi et al., with the simplest model, demonstrates an ability to produce multiple epidemic waves. In addition, the largest simulated population is smaller than most U.S. states and tuned for a narrow geographic area, which limits the usefulness of its results for state-level and national policy-makers.

The EpiCast simulator

The smaller scales of the coupled contagion ABMs described above contrast sharply with that of HPC-scale simulators like EpiCast. EpiCast simulates disease spread on a massive artificial surrogate population of the full U.S. containing over 322 million agents^{11,16}. In the baseline version of EpiCast, this behavior relies on a combination of fixed exogenous responses based on demographics and top-down public health interventions taken as input. The population used in EpiCast is built using UrbanPop²⁸, a framework which uses American Community Survey (ACS) Data to create high-fidelity synthetic populations. Compared to similar frameworks, UrbanPop captures more of the heterogeneity in real-world populations, producing cross-sectional representations of agents across sociodemographic, economic, educational, and mobility characteristics. UrbanPop assigns each member of the generated population to a specific household in a specific census block group, in a manner which ensures the aggregate demographics of the generated block group correspond precisely to those observed in census data. In EpiCast, these households form the basis for dividing the population into communities, within which agents have varying contact intensities defined by whether they share a household, neighborhood, workplace, school, etc. These contact intensities determine the pairwise likelihood that disease transmission occurs between two agents within the same community. Agents then commute between a home and a work or school community each day in a manner consistent with Department of Transportation commute flow data, and occasionally engage in long distance travel to geographically distant communities based on flight data.

Methods

In designing our model, we seek to adapt an existing coupled contagion ODE model to function within EpiCast, which requires accounting for all of the disease states present in the original code. We first construct a SIR-based ODE-model as a simple baseline, then incrementally incorporate new disease states until we have accounted for all those used by EpiCast. Once this design is complete, we then modify the model to function in an individualized context. In the process, we introduce both short- and medium-distance mechanisms for fear spread and two types of behavioral responses to fear with an eye toward reproducing the twin epidemic waves found in the ODE models.

Coupled contagion ODE models

We take a streamlined version of Epstein et al.'s model of coupled fear and disease spread²¹ as our starting point due to its flexibility and ability to produce multiple epidemic waves. We then incrementally add the asymptomatic, presymptomatic and exposed disease states present in EpiCast to this model, in order to ensure that the dynamics found in the simpler models are not disrupted by the addition of these disease states. The final model uses all disease states present in EpiCast. We refer to this series of models using the Cartesian product of the two types of states. For example, we describe a standard *SIR* model paired with a fear model with Neutral and Fearful states as an *SIR* \times *NF* coupled model. Note that these models are not intended to be directly compared with the extension of EpiCast described in “[Coupled contagion agent-based model](#)”. We instead use them to inform the high-level design of the model implemented in EpiCast, with a focus on how to couple the dynamics of disease, fear, and behavior.

We designed the following series of models:

1. *SIR* \times *NF*, the closest to Epstein et al.'s original model. The flow rate from neutral to fearful compartments is slower by a factor of ρ_f from the Recovered state. $\rho_f = 0$ corresponds to the Epstein model.
2. *SI_sI_aR_sR_a* \times *NF*, distinguishing between symptomatic and asymptomatic Infectious and Recovered states. Of the disease states, only *R_s* has a reduced flow rate from neutral to fearful compartments and contributes to a reduction in fearful agents. This is because agents are only aware of symptomatic infections, given that we do not currently model testing.

3. $SEPI_s I_a R_s R_a \times NF$, adds an Exposed state to represent the incubation period and a Presymptomatic state to allow for disease spread prior to the potential onset of symptoms. This is the closest to the model implemented in EpiCast.

For brevity, we present only the $SEPI_s I_a R_s R_a \times NF$ system in full. The $SEPI_s I_a R_s R_a \times NF$ flow diagram is shown in Supplementary Fig. S1 and the full system of equations is given in Supplementary Section S1. The other models are special cases of this system and can be derived by setting the the compartments not present in a given model to 0.

Coupled contagion agent-based model

The EpiCast coupled contagion model can be thought of as a translation of the $SEPI_s I_a R_s R_a \times NF$ model into the individualized context of EpiCast. This requires reworking aspects of the basic functionality of the original EpiCast simulator to facilitate fear spread and tie an individual's behavior to their fear state.

Modeling fear spread

We implement two methods of fear spread in EpiCast: one which operates based on an agent's local neighborhood of contacts in a manner similar to disease spread, and a second that approximates the effect of traditional newspaper publishers and broadcast media. For simplicity, we do not distinguish between these two forms of media, and refer to media outlets of either type as "broadcasters" below when describing the model.

In order to implement local fear spread, we took advantage of EpiCast's preexisting methodology for computing disease spread, refactoring the code to allow for fear spread using different per-contact transmission probabilities. Fear is spread by symptomatic and fearful agents, and countered by agents who know they have recovered from an infection and have a neutral fear state (i.e. agents in disease state R_s and fear state N).

For non-local fear spread, we created broadcasters capable of taking one of three positions on the disease. They may either: (1) spread fear, (2) counter fear spread, or (3) remain neutral. In order to site these broadcasters, we begin by examining the occupations of agents in each community, which are identified by 2017 North American Industry Classification System (NAICS) codes in EpiCast. If we identify at least one agent who works in publishing (NAICS code 511) or broadcasting (NAICS code 515) in a given community we locate a media outlet there and assign all agents within the relevant industry to work there. Each broadcaster is able to potentially influence anyone within the census tract in which they are located, and starts by taking a neutral position.

Once initialized, a broadcaster's position is determined by a combination of the fear states of its employees and the level of disease spread in the broadcast area. When the proportion of employees at a broadcaster who are fearful of the disease crosses a certain threshold, p_{bc_start} , the broadcaster begins spreading fear. We then record the current infection rate, the number of new infections in the broadcast area as a proportion of its population, as \dot{E}_0 . In future timesteps, the broadcaster determines its position by comparing the current infection rate, \dot{E}_t , against \dot{E}_0 as follows:

1. If $\dot{E}_t < p_{bc_counter} \cdot \dot{E}_0$, the broadcaster begins countering fear spread.
2. Otherwise, if $\dot{E}_t < p_{bc_neutral} \cdot \dot{E}_0$, the broadcaster takes a neutral position.
3. Otherwise, the broadcaster spreads fear of the disease.

where $p_{bc_counter} < p_{bc_neutral}$.

For each agent in the broadcast area of a given broadcaster, they have a fixed chance, p_{bc} , of being influenced by that broadcaster in a given timestep if the broadcast is not neutral and their fear state does not correspond to the broadcaster's position.

Modeling behavior

We modeled two different ways in which agents may change their behavior when they are fearful: (1) taking protective measures (such as masking and hand-washing) while maintaining their normal routine, and (2) withdrawing from their normal routine and remaining at home. We model this first measure by reducing the susceptibility of agents in state $< S, F >$, similar to how this is handled by Epstein et al.²² and in the ODE model.

In the second case, we modify an agent's routine so that they remain in their home community, neither commuting to their work community nor engaging in long distance travel, and only interacting with other members of their household. If currently traveling, they immediately return to their home community. Note that this is different from a complete quarantine, as they can still transmit the disease or their fear to their household. If agents are fearful, they can either withdraw purely due to fear with a fixed probability p_{fear} for a given duration t_{with} , or due to having a symptomatic infection (i.e. being in state $< I_s, F >$). Note that any agent with a symptomatic infection will withdraw if their symptoms are severe enough for them to be hospitalized, regardless of their fear state.

Experimental setup

First, we sought to examine how different sets of disease compartments impact outbreak dynamics (RQ 1). We find numerical solutions to each ODE model through utilizing an initial value problem (IVP) solver – the `solve_ivp` function from the Scipy integrate package (v1.14.0) – under different parameter choices. In the process, we identify scenarios under which each of our ODE models produce multiple epidemic waves, providing a comparison of when the same parameter values produce substantially different outcomes between models. Parameter values are shown in Supplementary Table S1.

Second, we sought to compare the outbreak dynamics resulting from several different sets of fear-induced behaviors (RQ 2). We run a range of scenarios in EpiCast where agents withdraw from everyday activities under certain conditions (most in part due to fear) and have a reduced susceptibility to infection due to protective actions. This includes withdrawals by agents who are (1) hospitalized (hospitalized withdrawals), (2) symptomatic and fearful (sick withdrawals), and (3) only fearful (fear-only withdrawals). These scenarios also compare the impact of purely local (based on physical contact with agents in the same community) and non-local (based on broadcasters operating within a given census tract) fear spread modes. For these scenarios, we use a synthetic population of about 322 million agents representing the 48 contiguous U.S. states and Washington DC. We seeded index cases on a per-county basis using estimates of active cases on 25 March 2020 derived from data from The New York Times, based on reports from state and local health agencies. Index cases were seeded on each of the first five days of the simulation, for an average of 856 thousand seed cases per day (about 0.265% of the population). Parameter values are shown in Supplementary Table S2.

Finally, we investigated how a couple key model parameters impact the dynamics produced by the behavioral model implemented in EpiCast (RQ 3). We conduct a series of sensitivity analyses comparing EpiCast outputs for a range of values of p_{fear} and σ_f , which determine the strength of an agent's behavioral response to fear of the disease. We draw these values from the range $[0, 1]$, trying five values for each parameter for a total of 25 combinations. We run each combination twice—once with only local fear spread and once with both local and broadcaster-based fear spread. These experiments are conducted on the contiguous U.S. dataset described above. Choices for parameters other than p_{fear} and σ_f correspond to Scenarios (c) and (d) from Supplementary Table S2 (or, equivalently, Scenarios (e) and (f)) without and with broadcasters, respectively. In order to compute the number of epidemic waves programmatically, we first compute a centered 7-day rolling maximum of new case counts—so as to avoid identifying peaks caused solely by weekly variation—then call the `find_peaks` function from the Scipy signal package (v1.14.0) in Python with a prominence of 0.0001. This value indicates that new cases should fall by at least 1 new infection per 10,000 people before rising again if we are to consider a given local maximum to be the peak of a given epidemic wave. We find this value avoids spurious peaks (i.e. those in the long tail of some outbreaks) while still capturing those most visually prominent in the data.

Results

We describe the results of the three main sets of experiments described above.

Comparing ODE models

For the ODE comparison experiments, Fig. 1 shows the solutions corresponding to the parameter choices that either failed to produce a second epidemic wave within 360 days (top) or produced such a wave (bottom). For the $SIR \times NF$ model, we observe that changing the likelihood that those who recovered from a symptomatic infection become newly fearful of the disease, ρ_f , has a dramatic impact on whether or not a second wave occurs. With $\rho_f = 1$ the percentages of the population in the susceptible or fearful states largely plateau around 80% and 60%, respectively, after 360 days, as shown in Fig. 1a. In contrast, these percentages plateau at around 35% and 0%, respectively, when we use $\rho_f = 0$, as shown in Fig. 1d. In addition, a second wave only occurs in the latter case, peaking after about 150 days.

When we introduce asymptomatic disease states, as in the $SI_sI_aR_sR_a \times NF$ model, the situation changes further. Keeping all parameters used in the $\rho_f = 1$ run the same with this model results in 80% of the population still susceptible and 40% fearful after 360 days, and no second epidemic wave, as shown in Fig. 1b. Changing the susceptibility of fearful agents, σ_f , from 25% to 35%, however, is sufficient to produce a second wave in this model, albeit a smaller one which peaks much later at around 275 days, as shown in Fig. 1e. Under this scenario, about half of the population is still susceptible after 360 days, although this proportion is still decreasing. The fearful proportion, while approaching zero by 360 days, does so about 160 days later than in the two wave scenario without asymptomatic infections (Fig. 1d).

Introducing an exposed state and a pre-symptomatic infectious state into the model produces the $SEPI_sI_aR_sR_a \times NF$ model, and further changes the solution space. Using the same pair of parameter values as with the $SI_sI_aR_sR_a \times NF$ model results in faster disease spread and decay in fear levels. Figure 1c shows that about 70% and 20% of the population is susceptible and fearful, respectively, after 360 days with $\sigma_f = 0.25$ in this model. Additionally, while we do not see a second wave peak within 360 days here, we do see the start of such a wave and the lack of a plateau of fear levels that corresponds to a single-wave outbreak in the other models. When using $\sigma_f = 0.35$ with the $SEPI_sI_aR_sR_a \times NF$ model, Fig. 1f shows that fear levels reach zero and a second epidemic peak occurs, as with the $SI_sI_aR_sR_a \times NF$ model. However, both these events happen much sooner in the former model, although still about 100 days later than in the $SIR \times NF$ model.

Altogether, with respect to RQ 1 we find that adding asymptomatic states to the ODE models significantly delays the progression of multiple waves and reduces the rate of infections and fear decay as well as the total attack rate, while adding an exposed state has the opposite effect. Neither addition significantly changes the peak of fear levels. As for RQ 3, we find that changing the values of ρ_f and σ_f both showed the potential to change both whether or not we observe a second wave and the total attack rate.

Comparing selected scenarios in EpiCast

We now examine the results of the various scenarios evaluated for the coupled contagion model implemented in EpiCast, so as to explore RQ 2. Figure 2a showcases the differences in the count, height, and timing of epidemic peaks between scenarios. In particular, we see the highest peak in the scenario with only hospitalization-based withdrawals (hosp), as well as the highest total attack rate. This is followed by the scenario where we allow withdrawals based on having symptoms (hosp+sick) and pure-fear withdrawals as well (hosp+sick+fear), which also show a single peak. The scenario with the smallest peak (and also the lowest total attack rate) comes

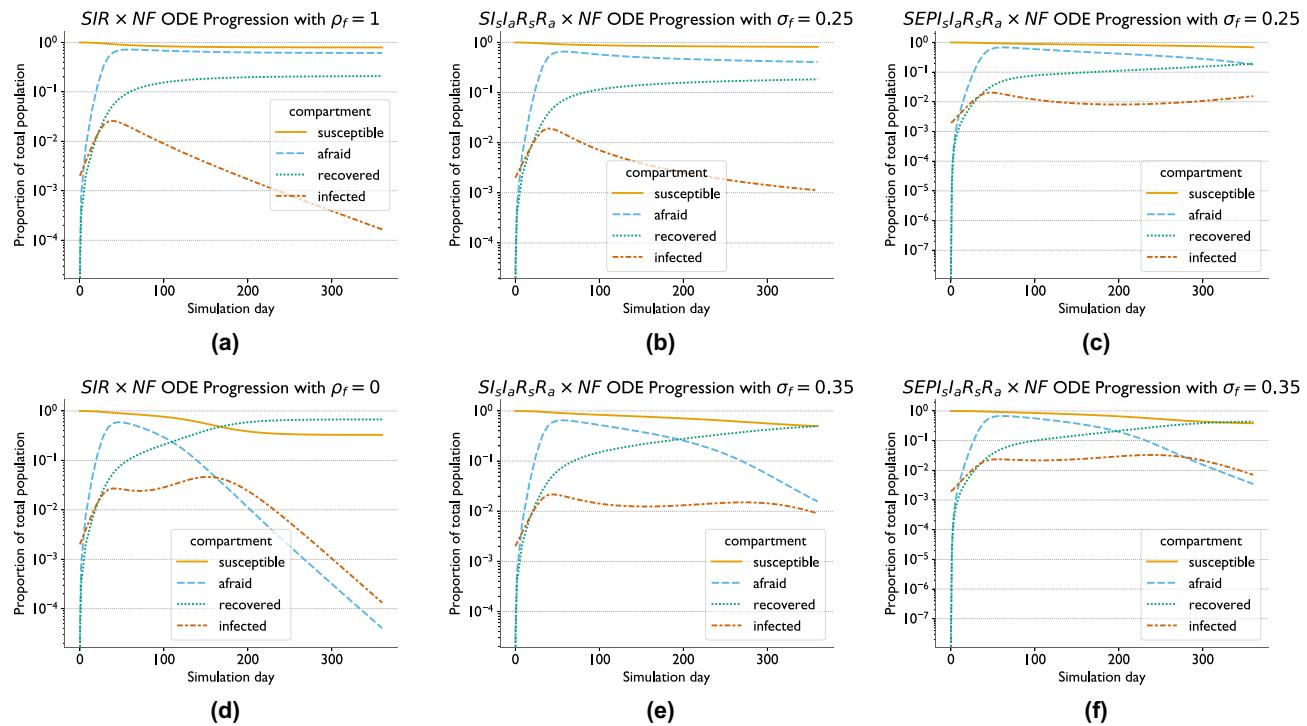


Fig. 1. ODE model comparison: We show outputs with attention to parameter choices which impact the emergence of multiple epidemic waves. On the top and bottom we show solutions corresponding to parameter choices which fail to produce, or produce, respectively, a second wave within 360 days. We show these results for models combining (1) a basic Susceptible Infectious Recovered model with Neutral and Fearful states ($SIR \times NF$, **a** and **d**, varying the relative likelihood an individual becomes fearful after recovering from a symptomatic infection, ρ_f), (2) separating symptomatic and asymptomatic infectious and recovered states ($SISI_aR_sR_a \times NF$, **b**) and introducing Exposed and Presymptomatic states (**e**, varying the susceptibility of fearful individuals to infection, σ_f), and $SEPI_sI_aR_sR_a \times NF$ (**c** and **f**, varying σ_f). For ease of comparison, we combine all infected (including E), recovered, and fearful compartments in each model into a single line in each of the plots above.

from using hospitalization-based and symptomatic withdrawals along with a susceptibility reduction from fear (hosp+sick+reduced_sus), which also has a single peak. The only scenarios wherein we observe multiple peaks are those where we add broadcasters as a fear spread mechanism with purely fear-based withdrawals (hosp+sick+fear+bc) and a fear-based susceptibility reduction (hosp+sick+reduced_sus+bc), with the former having higher peaks than the latter, and the second peak occurring at roughly the same time in both cases.

We observe that fear levels peak after 35 to 60 days at around 60-70% of the population in most cases, with the exception of the pure-fear withdrawal scenario without broadcasters (hosp+sick+fear), where fear levels peak after 85 days at around 38% of the population. In most cases, fear levels fall dramatically from this peak before the end of the simulation. The scenarios with reduced susceptibility due to fear both buck this trend, with fear levels only falling slightly by day 200 without broadcasters (hosp+sick+reduced_sus) and dropping slowly by around 15 percentage points in total with broadcasters (hosp+sick+reduced_sus+bc). We note that, in the absence of broadcasters, the reduced susceptibility case (hosp+sick+reduced_sus in Fig. 2b) displays the highest fear levels for most of the simulation and the pure-fear withdrawal case the lowest (hosp+sick+fear in Fig. 2b). In the presence of broadcasters, fear trends more closely resemble other runs. Fear levels fall significantly after their peak in the susceptibility reduction case (hosp+sick+reduced_sus+bc in Fig. 2b), when they remaining roughly constant in the absence of broadcasters. Conversely, fear levels rose much higher before falling in the pure-fear withdrawal case (hosp+sick+fear+bc in Fig. 2b), peaking at above 60% of the population rather than less than 40%.

In both scenarios with broadcasters, more broadcasters spread fear than countered it throughout the majority of the run, though in the pure-fear withdrawal case this changed shortly before the end of the run. Far more broadcasters were spreading fear in the susceptibility reduction scenario than in the pure-fear withdrawal scenario for the entire duration (compare hosp+sick+reduced_sus+bc and hosp+sick+fear+bc in Fig. 2c), while the reverse is true for fear countering for most of the duration of the run (Fig. 2d). Broadcasters were not used in the other scenarios.

Figure 3 highlights the geographical distribution of new cases for selected time steps and US states for two of the scenarios shown above. Figure 3a,b represent the results of the scenario with pure-fear withdrawals with only local fear spread, where Fig. 3c,d do so when fear spread through broadcasters is additionally allowed (hosp+sick+fear and hosp+sick+fear+bc, respectively, in Fig. 2). The time steps chosen represent peaks (days

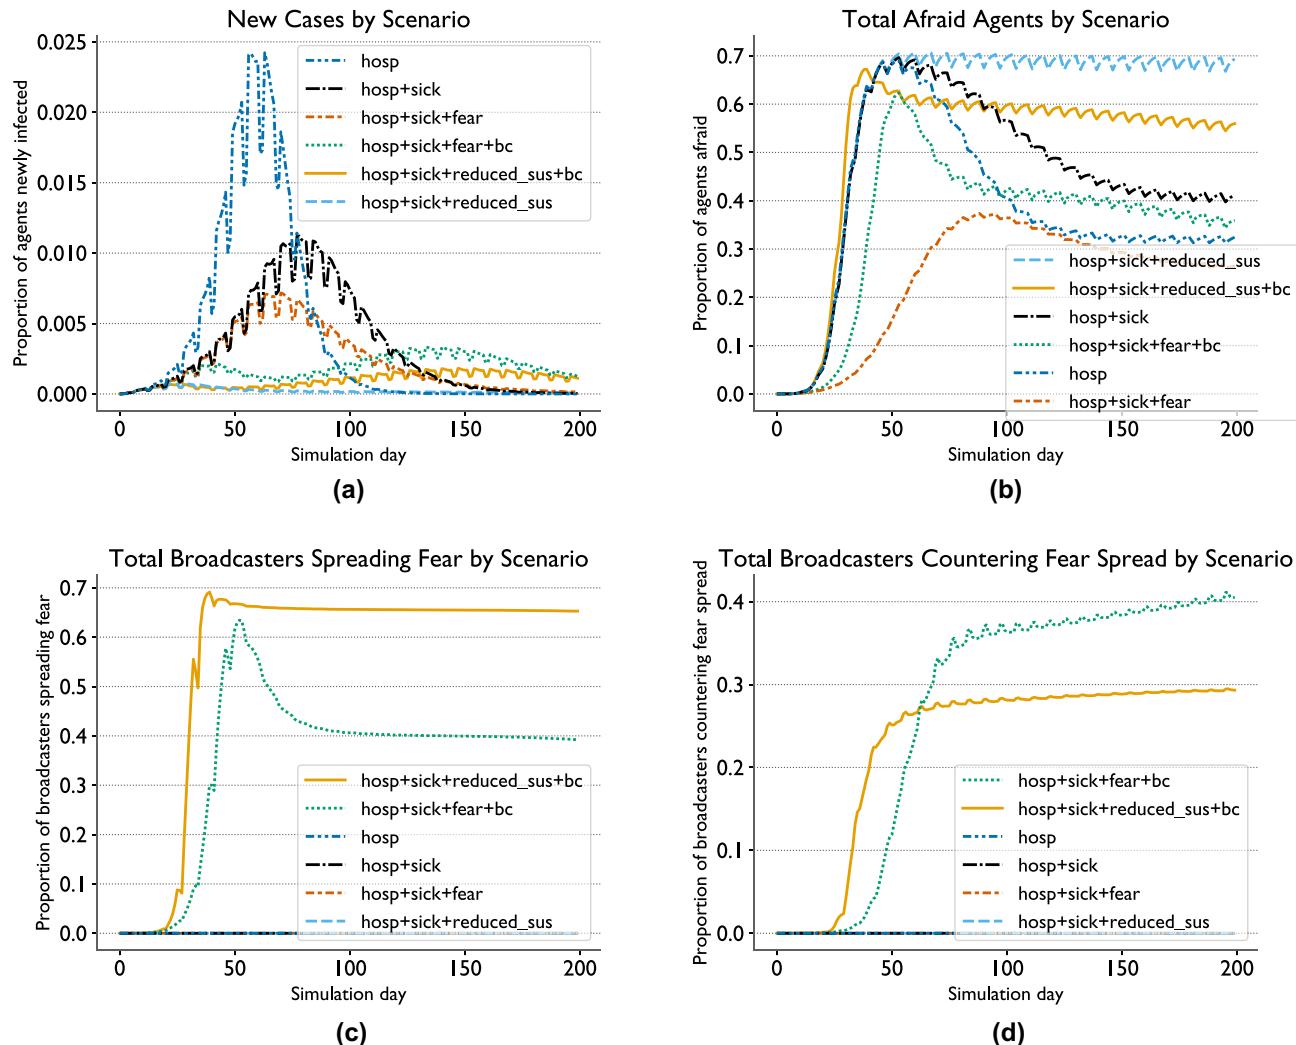


Fig. 2. EpiCast fear spread scenarios: new cases (a), total fear levels (b), total broadcasters spreading fear (c), and total broadcasters countering fear spread (d) for six different scenarios run in EpiCast. These scenarios can include agents withdrawing from their normal schedules due to being hospitalized (hosp), being fearful of the disease while having symptoms of it (sick), and fear of the disease without symptoms (fear), as well as non-local fear spread through broadcast media (bc), and fearful agents having a lower susceptibility to the disease due to taking protective actions such as masking (reduced_sus).

50, 100, and 130) and a trough (day 75) in the latter, multi-wave scenario. Particular US states were chosen to represent different epidemic trajectories.

For the former case, Fig. 3b shows how the state-level trajectories differ primarily based on the height of their peak, though Texas (TX) does generally lag behind the trends displayed by the other selected states by a handful of days, with the exception of Colorado (CO) in the declining phase of the outbreak.

In contrast, we see significant divergence between states in the broadcaster scenario, seen in Fig. 3d. South Dakota (SD) and Washington (WA) both have high initial peaks which transition into lower secondary peaks, with WA having an earlier first peak and consistently lower case counts throughout. Utah (UT), Colorado (CO), and Texas (TX), by contrast, each have low initial peaks and higher secondary peaks. For UT, there is very little decline in cases before the upward trend towards the second peak begins, with case counts peaking roughly 20–30 days before the other states. UT and CO reach similar heights to each other for each peak, while TX has a lower second peak, only slightly above WA.

EpiCast sensitivity analysis

In order to get a better sense of how varying the behavioral response to fear within EpiCast influences the overall trajectory of disease outbreaks, as per **RQ 3**, we ran a number of simulations varying two influential parameters. The first was the likelihood that an agent who is fearful of the disease—but not currently showing symptoms—withdraws from their regular schedule of activities purely due to fear, p_{fear} . The second was the scaling factor applied to a fearful agent's susceptibility to infection, σ_f , which represents the impact of fearful agents adopting

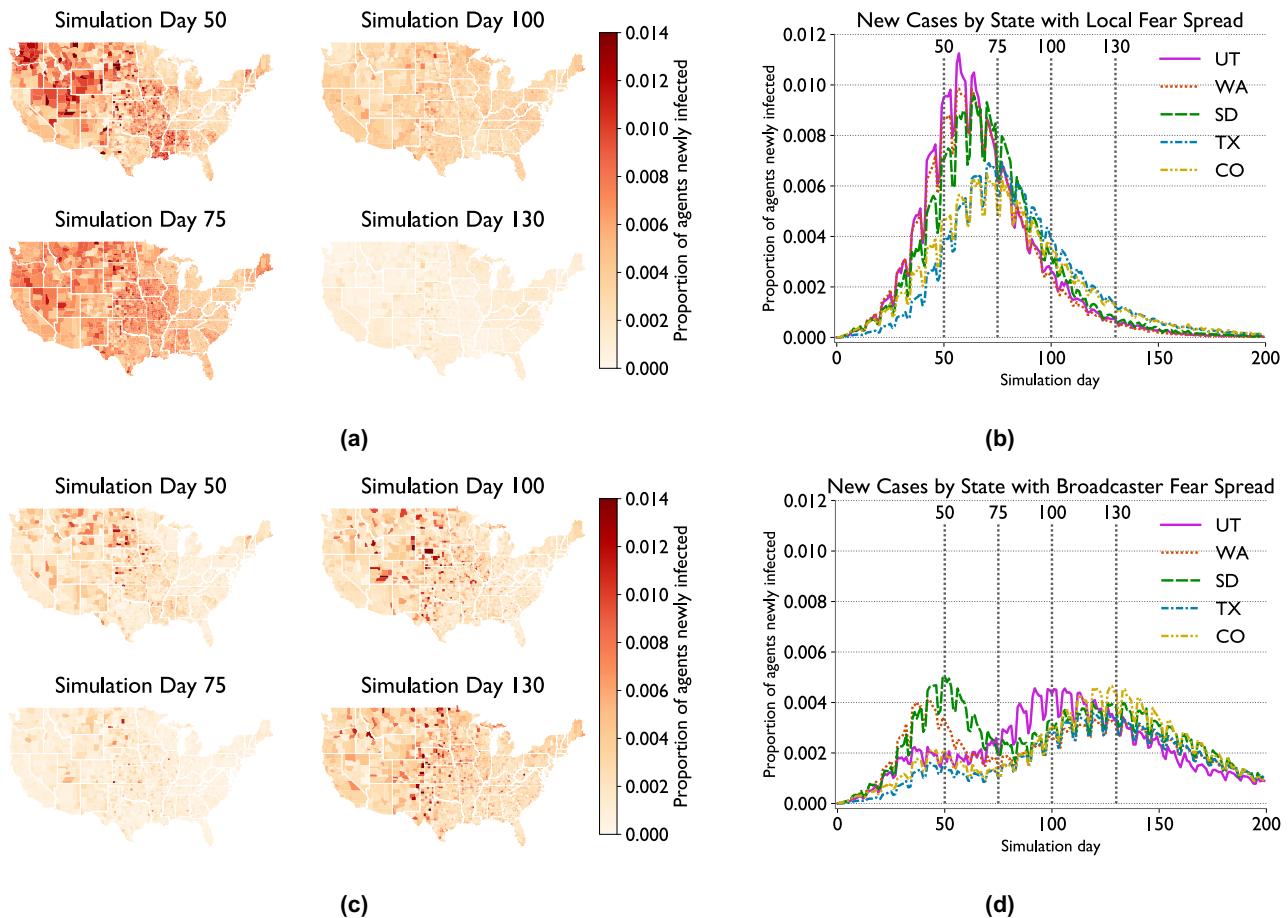


Fig. 3. Geographic distribution of outbreaks in two EpiCast fear spread scenarios: geographic distribution of new case counts for selected timestamps (left) and overall new case trends for selected states (right) for scenarios with pure-fear withdrawals without (upper) and with (lower) broadcaster-based fear spread.

protective behaviors that do not directly change their daily schedule, such as masking, maintaining social distance while in public, and hand-washing.

The results of conducting this parameter sweep with only local fear spread are shown in Fig. 4. Figure 4a displays the overall trajectories of each outbreak, with the proportion of new cases shown in orange, that of fearful agents shown in blue, and epidemic peaks indicated by the dashed vertical lines. Note that all parameter combinations produce a single epidemic wave. The simulation which comes the closest corresponds to $p_{\text{fear}} = 0$ and $\sigma_f = 0.5$. Further experiments, shown in Fig. 4b, revealed that a very narrow range of values – $p_{\text{fear}} = 0, 0.05, \sigma_f = 0.4$, and $p_{\text{fear}} = 0.1, \sigma_f = 0.45$ – produce multiple waves, but that a deviation by as little 0.05 in either parameter can be sufficient to return to the single wave scheme. Figure 4c shows the total attack rate of each simulated outbreak.

We observe the highest attack rate with $p_{\text{fear}} = 0$ and $\sigma_f = 1$, which represents a minimal behavioral response to fear (and corresponds to Scenario (b) from “Comparing selected scenarios in EpiCast”), followed by the cases with $p_{\text{fear}} = 1$ with lower values of σ_f . Notably, fear levels are much lower in the latter case. The lowest attack rates, in contrast, are observed when p_{fear} and σ_f are both small.

Higher values of p_{fear} correspond to fewer fearful agents and result in reduced variation in fear trajectories for different values of σ_f . For $p_{\text{fear}} \leq 0.5$, we observe that fear levels rapidly plateau for $\sigma_f \leq 0.25$ and rise to a peak before falling to varying degrees for $\sigma_f \geq 0.5$. This drop in fear levels is more pronounced for larger values of σ_f , when the overall attack rate is higher.

The results of conducting this parameter sweep with both local and broadcaster-based fear spread are shown in Fig. 5. The overall trajectories of each outbreak are shown in Fig. 5a with the dashed vertical lines indicating where peaks occur, and the orange and blue lines indicating proportion of new cases and fearful agents, respectively. Figure 5b highlights which combinations of parameters produce one and two epidemic waves. All single wave scenarios occur with low values of p_{fear} , and either a high or low value of σ_f (i.e. the upper right corner of Fig. 5b and the top two cells in the left of that figure). All other parameter combinations produce two waves. These two waves are generally about the same height with $p_{\text{fear}} = 1$ but the second is both taller and longer than the first in most other cases.

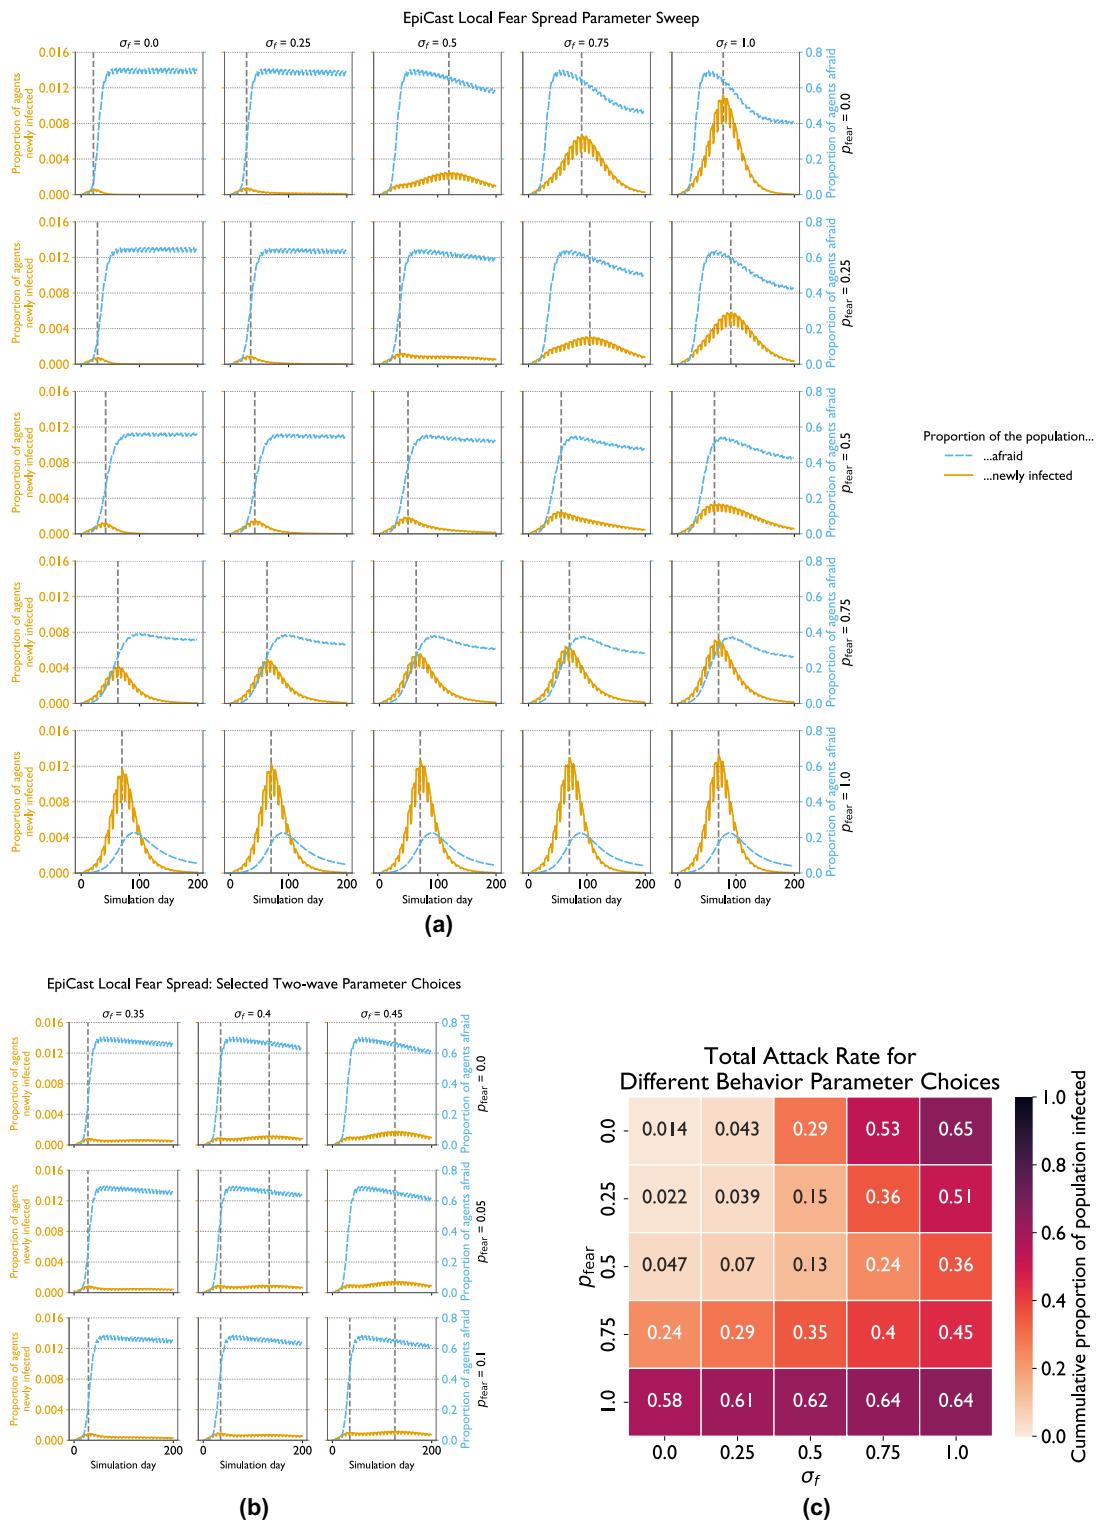


Fig. 4. EpiCast local fear spread sensitivity analysis: New cases by day (with peaks indicated by dashed vertical line) for (a) a coarse-grained parameter search and (b) a finer-grained search highlighting cases with multiple waves, and (c) total attack rate for various rates of purely fear-based withdrawals (p_{fear}) and levels of susceptibility scaling due to fear (σ_f) with only local fear spread. This parameter search compares different levels of susceptibility to infection for fearful agents (σ_f) and rates of withdrawal by fearful agents (p_{fear}).

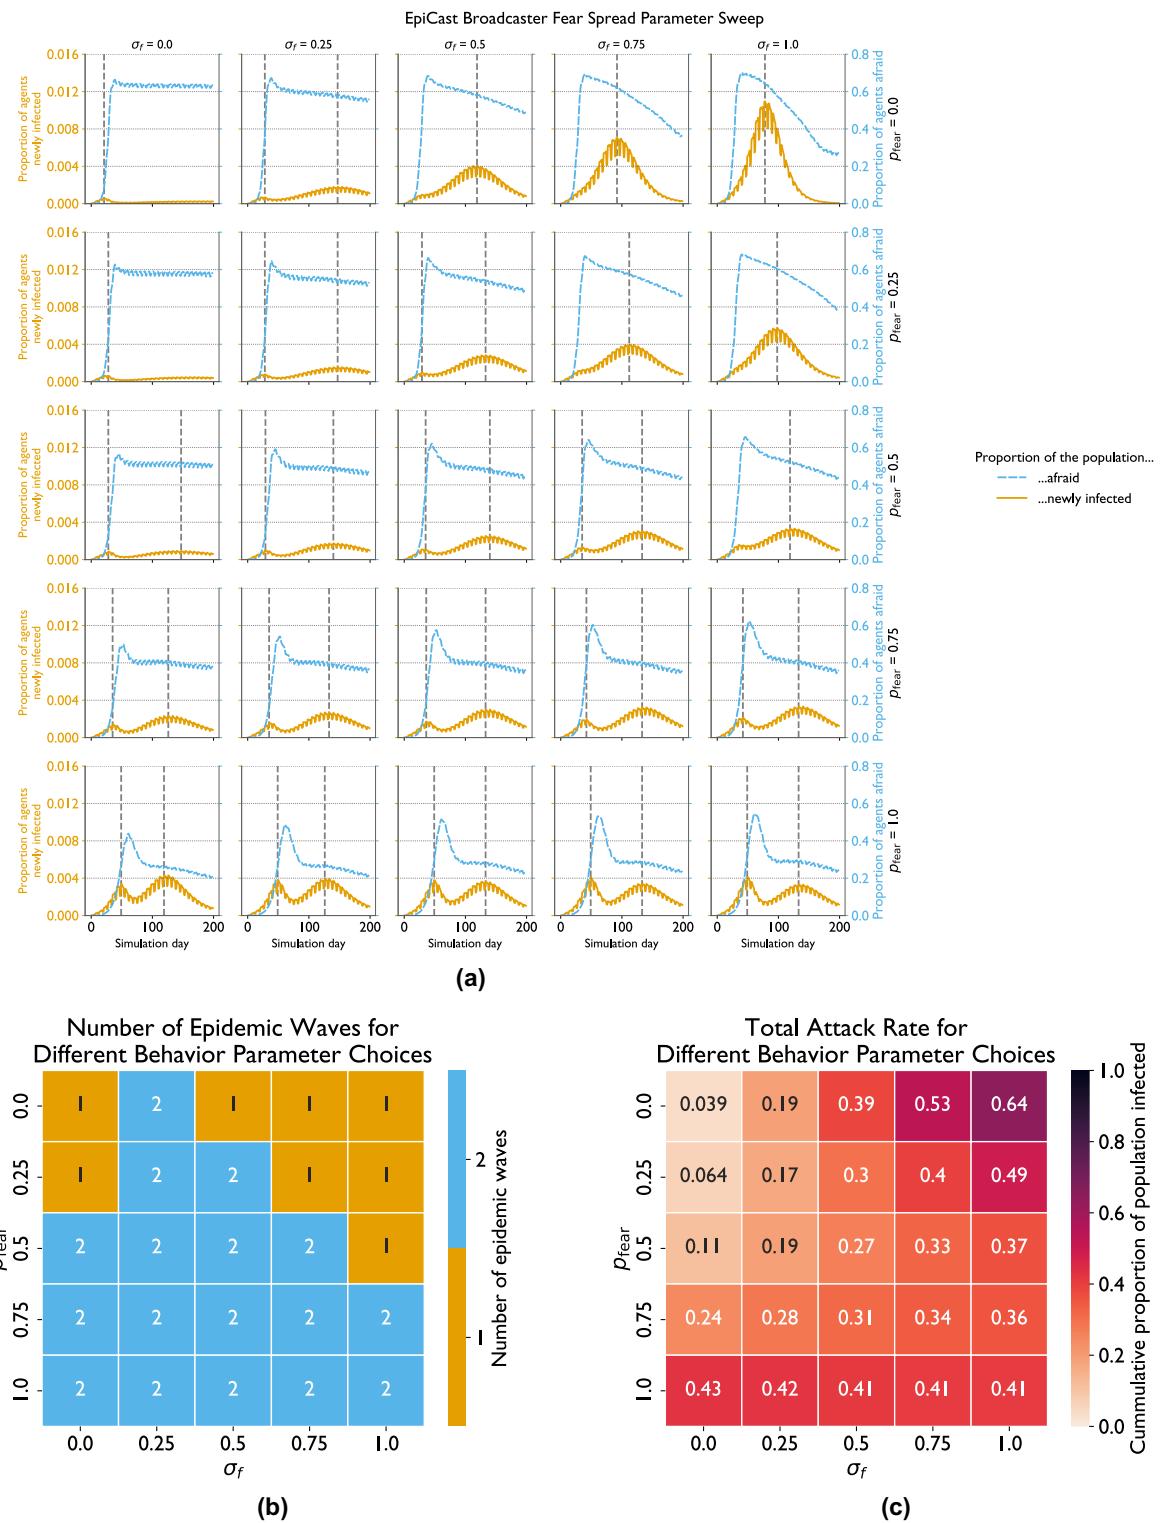


Fig. 5. EpiCast broadcaster fear spread sensitivity analysis: (a) new cases by day (with peaks indicated by dashed vertical line), (b) number of epidemic waves, and (c) total attack rate for various rates of purely fear-based withdrawals (p_{fear}) and levels of susceptibility scaling due to fear (σ_f) with both local and broadcaster-based fear spread. This parameter search compares different levels of susceptibility to infection for fearful agents (σ_f) and rates of withdrawal by fearful agents (p_{fear}).

The same parameter combinations that produce a single wave also produce the most extreme attack rates, as highlighted by Fig. 5c. The upper right corner (with high values of σ_f and low values of p_{fear}) contains cells with the highest total attack rates, while the single wave runs in the upper left corner exhibit the lowest. Compared to the corresponding simulations without broadcaster-based fear spread, those with broadcaster-spread fear tend to show more moderate attack rates. For example, attack rates fall from 58% – 64% (bottom row of Fig. 4c) to 43% – 41% (bottom row of Fig. 5c) when broadcasters are introduced with $p_{\text{fear}} = 1.0$. Conversely, attack rates largely rise in the top half of the first two columns of the figure, with for instance the attack rate with $p_{\text{fear}} = 0, \sigma_f = 0$ increasing from 1.4% to 3.9%, and that with $p_{\text{fear}} = 0.25, \sigma_f = 0.25$ rising from 7% to 19%.

Fear levels generally follow the same basic trajectory, rising to an initial peak before declining steeply for a short period, followed by a secondary slow decline. The transition between these two periods of decline generally corresponds to when cases begin rising for a second epidemic wave. The main exceptions are found in the top left and right corners of Fig. 5a, which show a plateau after a very short decline from the initial peak in the former case and a single accelerating decline for two cells in the latter. These mostly correspond to the cells with a single epidemic wave. The height of the initial peaks decreases as p_{fear} increases and fear levels fall farther when more cases occur.

Discussion

Examining the fear levels in each scenario, shown in Fig. 2b, offers some insights into why we might observe multiple peaks primarily in cases with non-local fear spread via broadcasters. In cases where fear reduces agent susceptibility to infection, in the absence of fear spread through broadcasters (hosp+sick+reduced_sus) fear levels never fall far below their peak, at around 50 days. This is likely due to infection levels never rising to the point where there is a sufficient number of agents in the R_s state to cause fearful agents to lose their fear. This is in contrast to what occurs with fear spread through broadcasters (hosp+sick+reduced_sus+bc), where fear levels quickly fall after reaching their peak. This suggests that the influence of broadcasters is sufficient to lower fear levels to the point where the disease can start spreading again, causing a second wave. Roughly the opposite appears to occur in the cases where we allow purely fear-based withdrawals; without the influence of broadcasters (hosp+sick+fear), fear levels remain lower than in any other scenario for the duration of the run. This is likely because withdrawals reduce *both* fear and disease spread in this mode and there is a high chance that any given fearful person will withdraw. Conversely, with the influence of broadcasters (hosp+sick+fear+bc) fear levels rise much higher before falling, though they remain lower than in all other scenarios except that with only hospitalization-based withdrawals. This suggests broadcasters allow fear to spread even when most fearful agents remain withdrawn.

These explanations are supported by examining the number of broadcasters spreading (see Fig. 2c) and countering fear spread (see Fig. 2d) over time. The period from days 25–50 where fear levels peak and then begin to fall in the reduced susceptibility scenario with broadcasters is also the period when we see the number of broadcasters countering fear spread dramatically increase, before reaching a much slower sustained growth rate. This slower growth period corresponds to a slow linear decrease in the number of fearful agents. The number of fear spreading broadcasters follows a similar trend to fear levels in this case, albeit with a smaller falloff after the peak. Similarly, we see that overall fear levels and the number of fear-spreading broadcasters show similar patterns in the pure-fear withdrawal case, with the dramatic reduction in fearful agents between days 50 and 100 matched by decline in fear-spreading broadcasters of similar magnitude.

In addition, the very narrow range of parameter values for which we do observe multiple waves with purely local fear spread appears to fall in the midst of a phase transition between a regime of small, short outbreaks and large, long outbreaks as represented by the cell $p_{\text{fear}} = 0, \sigma_f = 0.25$, and the cell $p_{\text{fear}} = 0, \sigma_f = 0.75$ in Fig. 4a, respectively. The former regime corresponds to lower values of σ_f which represent a strong behavioral response to fear, implying that fear-induced behavior reduces disease spread sufficiently to forestall a larger outbreak here. The latter regime, by contrast, represents minimal behavioral responses to fear, as seen by the fact it centers on $p_{\text{fear}} = 0, \sigma_f = 1$, where fear has no impact on the behavior of agents without symptoms. These two regimes also appear with broadcaster-based fear spread, and correspond to the only cells that do not exhibit multiple waves in Fig. 5a.

Altogether, these observations suggest that we observe two waves for opposite reasons in these cases. With reduced susceptibility, fear levels generally do not fall far enough for disease transmission to pick back up and produce a second wave without broadcasters countering fear spread. Conversely, with pure-fear withdrawals, fear levels generally fail to rise high enough fast enough for the first wave to leave a large reservoir of susceptible agents without broadcasters increasing the rate of fear spread. In either case, non-local fear spread, as represented by broadcasters, appears to facilitate a transition from high initial fear levels to low final fear levels in the presence of fear-based behaviors that significantly reduce disease spread. We posit that this transition only occurs under a narrow range of conditions in the case of purely local fear spread as a result of fear both spreading and receding relatively slowly. In contrast, the faster rate at which fear levels rise and fall in the presence of broadcasters may explain why multiple epidemic waves are more commonly observed under these conditions. In other words, in order to produce multiple epidemic waves in this model, fear must first spread significantly faster than the disease and then recede faster than it as well. Within this framing, it is clear why broadcasters facilitate the formation of multiple waves: information and the behavioral changes it engenders both spread and recede more rapidly when less tightly constrained by physical distance.

These results demonstrate that bottom-up drivers of behavior can shift outcomes in infectious disease simulations even in the absence of top-down interventions, like lockdown orders. Additionally, the difference in dynamics produced by adding non-local fear spread to our model further suggest that different structures of interaction between private actors may produce a range of different behaviors, even when the underlying

attitudes of the population remain unchanged. These findings complement observational work from the COVID-19 pandemic. One such study by Pan et al.²⁹ found that social distancing levels began rising prior to when lockdown orders were issued and fall before those orders were lifted in many U.S. states. Together, these results highlight the importance of securing broad buy-in from the public when implementing population-wide policy interventions in response to a pandemic.

The results we present in this paper are perhaps best understood as an exploration of the high-level implications of our model. Though the foundational elements of EpiCast are mature and have been calibrated extensively, there remain significant obstacles to doing so for the fear-based components of the model. While some survey data are available describing various behaviors throughout portions of the pandemic³⁰, it is much more difficult to tie behavior to specific influences or rationales. We thus confine ourselves to presenting the model itself and exploring the impact of varying parameter values, and leave calibration efforts to future work.

In addition, we acknowledge that the parameter sweep based sensitivity analysis study presented above is necessarily incomplete. While we explore the full possible *range* of parameter values for p_{fear} and σ_f , we do so at a relatively coarse resolution, so as to limit the total number of simulation runs. As evidenced by the multi-wave scenarios presented in Fig. 4b, this coarseness means that this study may miss dynamics that only emerge for a small range of parameter values. In addition, there are many simulation parameters beyond the two we chose to study. Extending the parameter sweep method to the other model parameters would be prohibitively computationally expensive, but in the absence of such an experiment it is unclear how varying those parameters might affect the dynamics in EpiCast.

Conclusion

In this work we present a series of models which capture the coupled spread of fear of a disease and the disease itself. We do so with several ODE models in order to demonstrate the impact of adding additional disease states to the model, representing asymptomatic, exposed, and pre-symptomatic disease states, and an ABM covering all disease states with both local and non-local mechanisms for fear spread implemented.

We find that the total attack rate is reduced by the addition of asymptomatic states and increased by the addition of a pre-symptomatic state in our ODE models, and the reverse is true for the rate at which an outbreak progresses.

In our ABM implementation in EpiCast, we primarily observe multiple epidemic waves when using our non-local fear spread mechanism, which represents broadcast media (regional information communication, including print, radio, and TV). We find that a combination of (1) high initial fear levels, (2) a strong behavioral response to fear, and (3) low fear levels later in the simulation produce multiple waves in our model. Thus, one potential explanation for why we only observe multiple waves with purely local fear spread for a very small range of parameter choices is that fear generally spreads (and recedes) much more slowly in the absence of broadcasters in our model.

Data availability

We have publicly released all output data from simulations described in this paper through Zenodo here: <https://zenodo.org/uploads/15706577>. The source code for the scripts used to solve the ODE models, analyze all simulation results, and produce the plots shown in this manuscript are publicly available here: <https://github.com/lanl/epicast-coupled>. The source code for EpiCast is available from Los Alamos National Laboratory, but restrictions apply to the availability of this code, which was used under license for the current study, and so is not publicly available. The code is, however, available from the corresponding authors upon reasonable request and with permission of Los Alamos National Laboratory.

Received: 1 July 2025; Accepted: 28 October 2025

Published online: 26 November 2025

References

- Chernozhukov, V., Kasahara, H. & Schrimpf, P. Causal impact of masks, policies, behavior on early covid-19 pandemic in the us. *J. Econ.* **220**, 23–62 (2021).
- Navas Thorakkatt, M., Farhin, S. & Khan, A. A. Forecasting the trends of covid-19 and causal impact of vaccines using bayesian structural time series and arima. *Ann. Data Sci.* **9**, 1025–1047 (2022).
- Toharudin, T. et al. National vaccination and local intervention impacts on covid-19 cases. *Sustainability* **13**, 8282 (2021).
- Suthar, A. B. et al. Public health impact of covid-19 vaccines in the us: observational study. *BMJ377* (2022).
- Cramer, E. Y. et al. Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the United States. *Proc. Natl. Acad. Sci.* **119**, e2113561119 <https://doi.org/10.1073/pnas.2113561119> (2022).
- Nixon, K. et al. An evaluation of prospective covid-19 modelling studies in the usa: from data to science translation. *Lancet Digital Health* **4**, e738–e747 (2022).
- Bauch, C., d’Onofrio, A. & Manfredi, P. Behavioral epidemiology of infectious diseases: an overview. *Modeling the interplay between human behavior and the spread of infectious diseases* 1–19 (2013).
- Hamilton, A. et al. Incorporating endogenous human behavior in models of covid-19 transmission: A systematic scoping review. *Dial. Health* 100179 (2024).
- Schlüter, M. et al. Unraveling complex causal processes that affect sustainability requires more integration between empirical and modeling approaches. *Proc. Natl. Acad. Sci.* **120**, e2215676120. <https://doi.org/10.1073/pnas.2215676120> (2023).
- Murakami, T., Sakuragi, S., Deguchi, H. & Nakata, M. Agent-based model using gps analysis for infection spread and inhibition mechanism of sars-cov-2 in tokyo. *Sci. Rep.* **12**, 20896 (2022).
- Germann, T. C., Kadau, K., Longini, I. M. Jr. & Macken, C. A. Mitigation strategies for pandemic influenza in the united states. *Proc. Natl. Acad. Sci.* **103**, 5935–5940 (2006).
- Kersting, M., Bossert, A., Sørensen, L., Wacker, B. & Schlüter, J. C. Predicting effectiveness of countermeasures during the covid-19 outbreak in south africa using agent-based simulation. *Hum. Soc. Sci. Commun.* **8**, 1–15 (2021).

13. Bosman, M. et al. An agent based simulation of covid-19 history in catalonia using extensive real datasets. *Sci. Rep.* **14**, 31858 (2024).
14. Chen, J. et al. Role of heterogeneity: National scale data-driven agent-based modeling for the us covid-19 scenario modeling hub. *Epidemics* **48**, 100779 (2024).
15. Mao, L. Modeling triple-diffusions of infectious diseases, information, and preventive behaviors through a metropolitan social network—an agent-based simulation. *Appl. Geogr.* **50**, 31–39 (2014).
16. Alexander, P. C. et al. Epicast 2.0: A large-scale, demographically detailed, agent-based model for simulating respiratory pathogen spread in the United States, <https://doi.org/10.48550/arXiv.2504.03604> (2025). [arXiv:2504.03604](https://arxiv.org/abs/2504.03604).
17. Rahmandad, H., Xu, R. & Ghaffarzadegan, N. Enhancing long-term forecasting: Learning from covid-19 models. *PLoS Comput. Biol.* **18**, e1010100 (2022).
18. Tovissodé, C. F. & Baumgaertner, B. Heterogeneous risk tolerance, in-groups, and epidemic waves. *Front. Appl. Math. Stat.* **10**, 1360001 (2024).
19. Kassa, S. M. & Ouhinou, A. Epidemiological models with prevalence dependent endogenous self-protection measure. *Math. Biosci.* **229**, 41–49 (2011).
20. Huang, Y. & Zhu, Q. Game-theoretic frameworks for epidemic spreading and human decision-making: A review. *Dyn. Games Appl.* **12**, 7–48 (2022).
21. Epstein, J. M., Parker, J., Cummings, D. & Hammond, R. A. Coupled contagion dynamics of fear and disease: mathematical and computational explorations. *PLoS ONE* **3**, e3955 (2008).
22. Epstein, J. M., Hatna, E. & Crodelle, J. Triple contagion: a two-fears epidemic model. *J. R. Soc. Interface* **18**, 20210186 (2021).
23. Poletti, P., Ajelli, M. & Merler, S. The effect of risk perception on the 2009 h1n1 pandemic influenza dynamics. *PLoS ONE* **6**, e16460 (2011).
24. Jovanović, R., Davidović, M., Lazović, I., Jovanović, M. & Jovašević-Stojanović, M. Modelling voluntary general population vaccination strategies during covid-19 outbreak: influence of disease prevalence. *Int. J. Environ. Res. Public Health* **18**, 6217 (2021).
25. Rajabi, A., Mantzaris, A. V., Mutlu, E. C. & Garibay, O. O. Investigating dynamics of covid-19 spread and containment with agent-based modeling. *Appl. Sci.* **11**, 5367 (2021).
26. Palomo-Briones, G. A., Siller, M. & Grignard, A. An agent-based model of the dual causality between individual and collective behaviors in an epidemic. *Comput. Biol. Med.* **141**, 104995 (2022).
27. Prieto Curiel, R. & González Ramírez, H. Vaccination strategies against covid-19 and the diffusion of anti-vaccination views. *Sci. Rep.* **11**, 6626 (2021).
28. Tuccillo, J. et al. Urbanpop: A spatial microsimulation framework for exploring demographic influences on human dynamics. *Appl. Geogr.* **151**, 102844 (2023).
29. Pan, Y. et al. Quantifying human mobility behaviour changes during the covid-19 outbreak in the united states. *Sci. Rep.* **10**, 20742 (2020).
30. Carnegie Mellon University, Delphia Group. COVID-19 Trends and Impact Survey (CTIS) Results. <https://delphi.cmu.edu/covidcast/survey-results/?date=20220625##symptoms>, last accessed Nov 8th, 2024 (2022).

Acknowledgements

Many thanks to Nidhi Parikh, Micaela Richter, Jeffery S. Keithley, and Thomas J. Harris for helpful discussions over the course of this work.

Author contributions

J.K., P.C.A, D.J.B., C.B., S.Y.D.V., and T.C.G. developed the conceptual framework and J.K. wrote the majority of the code for the behavioral models presented in this work. This code included the ODE models, and the extension to EpiCast, with P.C.A. contributing some code to the later. D.J.B. and C.B. provided suggestions regarding the development and evaluation of the ODE models. P.C.A. and T.C.G. provided insights into the original design and implementation of EpiCast which informed the implementation of the behavioral model. J.T. developed and provided access to UrbanPop, the synthetic population on which this study was conducted, and guided that use. J.K., A.B., S.Y.D.V., and T.C.G. conceived the experiments, and J.K. conducted them and analyzed the results. J.K. prepared the initial draft of the manuscript, along with all figures. A.B., S.Y.D.V., and T.C.G. all provided supervision. All authors reviewed and edited the manuscript.

Funding

This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and by cooperative agreement CDC-RFA-FT-23-0069 from the CDC's Center for Forecasting and Outbreak Analytics. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the Centers for Disease Control and Prevention.

This work was performed at Los Alamos National Laboratory (LANL), an equal opportunity employer, which is operated by Triad National Security, LLC, for the National Nuclear Security Administration (NNSA) of the U.S. Department of Energy (DOE) under contract #19FED1916814CKC. This work was also performed in part at Oak Ridge National Laboratory (ORNL) which is operated by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The findings and conclusions in this report are those of the authors and do not necessarily represent the official position of LANL. This research used resources provided by the Darwin testbed at LANL which is funded by the Computational Systems and Software Environments subprogram of LANL's Advanced Simulation and Computing program (NNSA/DOE). This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-26425-y>.

Correspondence and requests for materials should be addressed to J.K. or T.C.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025