

Interplay of epidemic spreading and vaccine uptake under complex social contagion

Alfonso de Miguel-Arribas*

*Institute for Biocomputation and Physics of Complex Systems (BIFI),
University of Zaragoza, 50018, Zaragoza, Spain. and
Zaragoza Logistics Center (ZLC), Av. de Ranillas 5, 50018, Zaragoza, Spain*

Alberto Aleta

*Institute for Biocomputation and Physics of Complex Systems (BIFI),
University of Zaragoza, 50018, Zaragoza, Spain. and
Department of Theoretical Physics, University of Zaragoza, 50018, Zaragoza, Spain.*

Yamir Moreno

*Institute for Biocomputation and Physics of Complex Systems (BIFI),
University of Zaragoza, 50018, Zaragoza, Spain. and
Department of Theoretical Physics, University of Zaragoza, 50018, Zaragoza, Spain. and
Centai Institute, Turin, Italy
(Dated: September 18, 2025)*

Modeling human behavior is essential to accurately predict epidemic spread, with behaviors like vaccine hesitancy complicating control efforts. While epidemic spread is often treated as simple contagion, vaccine uptake may follow complex contagion dynamics, where individuals' decisions depend on multiple social contacts. Recently, the concept of complex contagion has received strong theoretical underpinnings thanks to the generalization of spreading phenomena from pairwise to higher-order interactions. Although several potential applications have been suggested, examples of complex contagions motivated by real data remain scarce. Surveys on COVID-19 vaccine hesitancy in the US suggest that vaccination attitudes may indeed depend on the vaccination status of social peers, aligning with complex contagion principles. In this work, we examine the interactions between epidemic spread, vaccination, and vaccine uptake attitudes under complex contagion. Using the SIR model with a dynamic, threshold-based vaccination campaign, we simulate scenarios on an age-structured multilayer network informed by US contact data. Our results offer insights into the role of social dynamics in shaping vaccination behavior and epidemic outcomes.

I. INTRODUCTION

The integration of human behavior into the modeling of epidemic spreading has long drawn researchers' attention [1]. Human behavior fundamentally influences the spread of epidemics in various ways. These include how our societies, mindsets, and lifestyles are organized, how human settlements tend to favor crowding (both human and animal), as well as our modes of transportation, which facilitate the fast and global spread of otherwise local diseases. Additionally, our responses to epidemic outbreaks play an important role. Particularly in this latter aspect, modeling human behavior remains a formidable challenge due to its complex nature and the lack of real-world, quantitative data on behavioral changes in populations affected by epidemic outbreaks [2, 3].

Key aspects of human behavior that significantly shape epidemics are the fear of infection, information awareness, and behavioral attitudes towards control interventions [3–8]. In particular, vaccination provokes hesitant attitudes in some sectors of the population, diminishing its efficacy at the population level [9, 10]. For this reason, the dynamics of vaccination has long been a topic within

the epidemic modeling literature [11–15], and numerous works have attempted to incorporate genuine behavioral features into the vaccination process [4, 7, 15–17].

Despite the lack of a well-established unifying framework to couple epidemic spreading and social dynamics, one of the standard approaches to integrate human behavior into epidemiological models is game theory through the so-called vaccination games [18, 19]. Here, individuals decide whether to vaccinate or not depending on the assessment of the cost incurred by the options they face. This assessment although traditionally assumed to be perfectly rational and mostly isolated from the social environment, has been progressively extended to incorporate more realistic features that affect the human decision-making process [16, 18–20]. Another approach for integrating epidemic models and vaccination-related behavior, significantly less exploited than vaccination games, is that of opinion dynamics [21]. In this approach, individuals hold an opinion on a topic (such as vaccine uptake), which can change or be influenced through various mechanisms, including peer interaction and social and mass media exposure. Notable references include [22–28], where complex behaviors involving sudden transitions, bistability, or network segregation have been demonstrated.

Among the most notable models of opinion dynamics, lies the seminal Watts-Granovetter threshold model [29–

* alfonso.demiguel.arribas@gmail.com

31]. Originally aimed at modeling phenomena such as riots and other social movements, this model was later generalized to the study of information cascades, stock market crashes, and cascading failures in infrastructure networks. These collective phenomena are typically referred to as *complex contagions* [32]. Unlike simple contagions, where the transmission of infection occurs at a rate between individuals and a single infected neighbor is always sufficient to expose a susceptible node, in complex contagions, the behavior of an individual is conditional on the behavior of a fraction of their peers [32–34]. Thus, while biological processes such as epidemics may be seen as simple contagion processes, human behavior dynamics related to opinions and decisions on matters of self-protection and vaccine uptake might constitute complex contagions. Recent studies have indeed begun to explore the interaction of simple and complex contagions within the context of epidemics [6, 27, 35–40].

At the same time, a substantial body of theoretical work is being developed around the concept of complex contagion, through the generalization of spreading phenomena from pairwise to higher-order interactions, extending traditional networks to structures such as simplicial complexes and hypergraphs [41–43]. Currently, though several potential applications have been suggested [44], examples motivated by real data where complex contagion is at play are scarce. In this regard, we have found explicit evidence from COVID-19 vaccine hesitancy surveys [45], suggesting that people’s attitudes towards vaccine uptake could depend on the behavior of their social peers, perfectly aligning with the concept of a complex contagion process.

Motivated by these observations, in this work, we investigate the interplay between the spread of an epidemic on networks and the dynamics of opinion on vaccine uptake following a complex contagion interaction. We employ the standard SIR model, which we couple with a dynamic vaccination campaign. The progression of this campaign is, in turn, influenced by a threshold-based opinion dynamics process, effectively integrating models of simple and complex contagion. To offer a more realistic setting where the dynamics unfold, we use an age-based multilayer contact network substrate, generated from data-driven contact matrices in the United States [46]. We investigate the system behavior by first assuming that everyone has the same activation threshold to switch their opinion towards a pro-vaccine stance, and later we explore the heterogeneous landscape of vaccine attitudes across the US as provided by Lazer et al. surveys [45].

II. MODEL

A. Definition

Our model consists of coupled dynamics featuring a simple epidemic process, a vaccination campaign, and

opinion dynamics based on a complex contagion model, running on an age-layered contact network.

This multilayer network is generated from contact matrices constructed from age-based data-driven synthetic populations [46, 47] (see Methods, Section V A, for implementation details). In this construction, nodes in layer a represent individuals of age a , and their contacts are distributed both within their own layer and across other layers, following realistic contact patterns. Our geographical focus here is on the United States, both at the national level and on a state basis.

As a simple approximation to influenza-like illnesses (ILI), we take the conventional susceptible-infected-removed SIR model for the epidemic process:

$$S + \mathcal{I} \xrightarrow{\beta} \mathcal{I} + \mathcal{I}, \quad (1)$$

$$\mathcal{I} \xrightarrow{\mu} \mathcal{R}, \quad (2)$$

where β is the disease’s transmission rate, and μ is the recovery rate. Throughout this work, we fix the basic reproduction number at $R_0 = 1.5$, representing a moderate transmissibility scenario, and set $\mu = 0.2$, corresponding to an average infectious period of $1/\mu = 5$ days. These values are consistent with empirical studies and prior literature on influenza epidemiology [48–50]. Although the survey data informing opinion thresholds was collected during the COVID-19 pandemic, the epidemic model itself is not intended to reproduce COVID-19 dynamics, but rather to explore generic scenarios of mild, influenza-like illnesses. The transmission rate β is then adjusted following the relationship $R_0 \approx (\beta/\mu)\langle k \rangle$, where $\langle k \rangle$ denotes the average degree of the contact structure, which is approximately $\langle k \rangle = 12$ across all the states in the country (see Methods, Section V A, for details on the contact structure).

All susceptible individuals are eligible to an ongoing vaccination campaign, modeled also as a Poisson process with rate α :

$$S \xrightarrow{\alpha} \mathcal{V}. \quad (3)$$

This parameter α can be interpreted as the vaccination effort carried out by the public health authorities during the vaccination campaign. However, not all susceptible individuals will automatically proceed to vaccination, as this depends both on their individual attitude toward the control measure and on the ongoing progress of the vaccination campaign, thereby establishing a tight feedback loop between the two dynamical processes.

Inspired by surveys on vaccine hesitancy conducted at the beginning of the COVID-19 vaccination campaign in the US, we add a third dynamical process in the form of opinion dynamics under complex contagion. The attitude towards vaccination is binary: inactive (hesitant) \mathcal{H} or (pro-)active \mathcal{A} . The following transition is defined:

$$\mathcal{H} \rightarrow \mathcal{A}. \quad (4)$$

Thus, hesitant behavior is not necessarily anti-vaccine behavior and once an individual has decided to be vaccinated, they will not change their mind. The specific mechanism under which this change of opinion proceeds is through a threshold-based updating rule in the vein of the renowned Watts-Granovetter model [31]:

$$o_i(t+1) = \begin{cases} \mathcal{A} & \text{if } \frac{V_{\Omega_i}(t)}{k_i} \geq \theta_i, \\ \mathcal{H} & \text{otherwise.} \end{cases} \quad (5)$$

Here θ_i is the individual's i activation threshold, k_i is the number of first-neighbors of i , and $V_{\Omega_i}(t)$ is the total number of vaccinated individuals within these neighbors. This relationship directly connects the success of the ongoing vaccination campaign to the opinion on vaccination, setting a co-evolving feedback loop between vaccination and opinion dynamics. An alternative modeling choice could have been to base the updating rule on peers' declared vaccination stance (i.e., whether they intend to vaccinate, \mathcal{A} status, or not, \mathcal{H} status) rather than on their actual vaccination status. However, our approach deliberately couples the success of the vaccination campaign with the evolution of attitudes, allowing us to directly leverage survey data on observed vaccine uptake.

It is important to emphasize that only hesitant susceptible individuals participate in the opinion dynamics and only those who adopt a pro-vaccination attitude (\mathcal{A}) may proceed to vaccination. Infected and removed individuals do not reassess their attitudes towards vaccination and are excluded from the vaccination process altogether. This assumption reflects the limited preventive role of vaccination once individuals are already infected or have recovered and ensures that behavioral changes occur only among those who can still be protected.

Additionally, it must be noted that no age-dependence has been made explicit so far since all these processes and parameters are equal for any individual in the system. In our model, age enters implicitly through the specific contact patterns of the multilayer network. The reason for this is to focus directly on the global impact of the coupled simple and complex contagion dynamics, without confounding features.

Our approach is summarized in Figure 1. More details on the specifics of the data and modeling approach can be found in the Methods, section V.

B. Experiments

We run extensive discrete-time Monte Carlo simulations of the coupled dynamics on the multilayer structure.

The system behavior is characterized by these main observables: the number of active (convinced to vaccinate) agents $N_A(\infty)$, the prevalence $R(\infty)$, that is, the fraction of the population affected by the disease, and the vaccination coverage (VC) $V(\infty)$, the number of people that

end up taking the vaccine. All these quantities are referred to the equilibrium state, that is when the fraction of infected individuals in the system reaches zero. By normalizing these quantities by the system's size N , we represent them as $r(\infty) = R(\infty)/N$, $v(\infty) = V(\infty)/N$, and $n_A(\infty) = N_A(\infty)/N$, respectively. These quantities will be examined across varying control parameters such as the activation threshold θ , the initial fraction of active individuals $n_A(0)$, and the vaccination rate α . As a nuance, given that the outlook of $n_A(\infty)$ and $v(\infty)$ is very similar, and the fact that $n_A(\infty)$ is highly conditioned on $n_A(0)$, we define the following derived quantity:

$$\Delta n_A(\infty) \equiv \frac{n_A(\infty) - n_A(0)}{1 - n_A(0)}, \quad (6)$$

which can be interpreted as a measure of the normalized final size of the behavioral cascade triggered given the initial condition $n_A(0)$. In other words, $\Delta n_A(\infty)$ informs us about the change in vaccine support relative to the initial support $n_A(0)$. It must be noted that the distribution of this variable $n_A(0)$ is assigned uniformly at random, that is, the initial set of individuals with a proactive attitude towards vaccine uptake is independent of their age, degree, or opinion assortativity.

Finally, we consider vaccination rates ranging from $\alpha = 0.001$, which amounts to vaccinating 0.1% of the eligible population, to $\alpha = 0.05$. To put these numbers in perspective, the highest daily vaccination record in the United States, according to [51], was 3.15 million doses on April 11, 2021, when the total population that had received two doses was approximately 73 million, as reported by [52]. Given that the total US population is about 350 million, this represents a rate of roughly 0.012 on that record day (1.2% of the population). Nonetheless, actual vaccination rates are much more variable as compared to our simplified model of a constant rate. From a global perspective, during the COVID-19 pandemic, despite the urgent vaccination deployment efforts and increased disease awareness, the daily percentage of the population receiving a COVID-19 vaccine dose rarely exceeded 1%, equivalent to $\alpha = 0.01$, with occasional exceptions in countries like China and the UK. As reported in [53], these countries experienced peaks of 1.57% and 1.41% of their populations getting vaccinated in a single day, respectively. These figures are instrumental in contextualizing our vaccination rate scenarios.

III. RESULTS

A. Homogeneous thresholds

Initially, we study the system's behavior under a homogeneous threshold θ scenario (Figure 2). The parameter space related to the opinion dynamics process, $(\theta, n_A(0))$ is scrutinized in finer detail, and we select certain vaccination scenarios based on the rate α . As representative

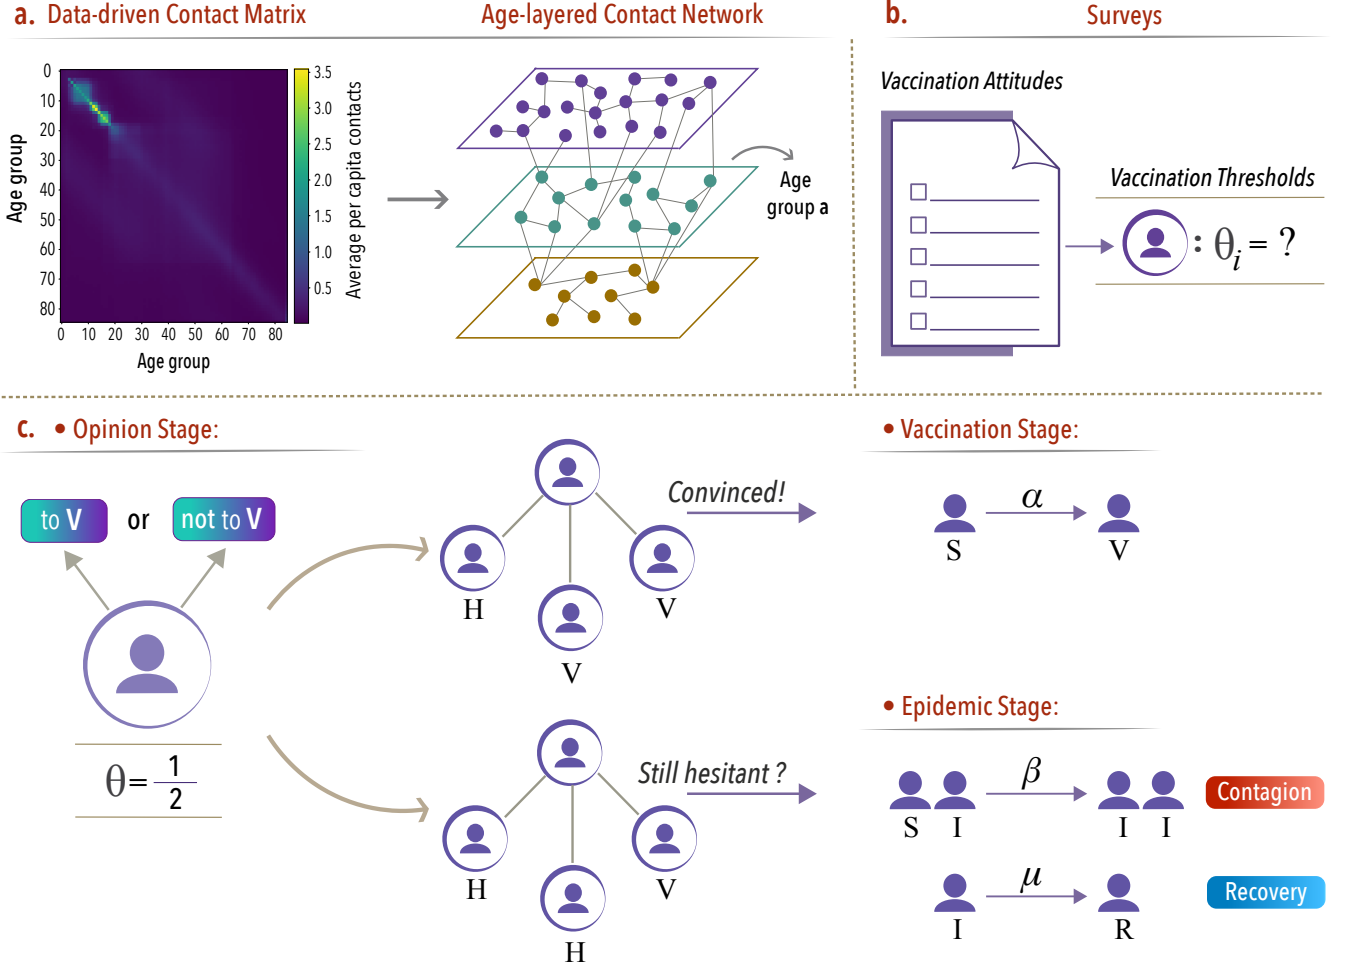


FIG. 1. **Modeling approach scheme.** (a) Contact matrices are constructed using data-driven synthetic populations [46] and are used to generate age-layered contact networks following the approach proposed in [47]. (b) Vaccination surveys conducted during the COVID-19 pandemic by Lazer et al. [45] inform the conversion of vaccination attitudes into numerical vaccination thresholds. Initially, individuals are assumed to have homogeneous vaccination thresholds. In a later stage of the work, these thresholds are refined using the data-informed values from the surveys. (c) With the contact structure established and vaccination thresholds assigned to individuals, the dynamics of the model are initiated. During the opinion stage, each hesitant susceptible individual decides whether to vaccinate based on their threshold and the vaccination status of their neighbors. If the proportion of vaccinated neighbors exceeds the threshold θ , the individual becomes vaccinated with probability α . Otherwise, they remain unvaccinated and may proceed to the epidemic stage governed by a standard SIR model. Only susceptible individuals are subject to opinion dynamics and the vaccination process; infected and removed individuals do not reassess their attitudes or undergo vaccination.

enough, the US contact patterns data at the national level are utilized to generate the underlying multilayer contact network.

Four different vaccination scenarios with substantially different vaccination efforts are considered, ranging from $\alpha = 0.001$ (Panels a1-c1), $\alpha = 0.005$ (Panels a2-c2), $\alpha = 0.01$ (Panels a3-c3), to $\alpha = 0.05$ (Panels a4-c4). As framed by the prior discussion on COVID-19 vaccination rates, the scenarios proposed here lie from a very low effort, $\alpha = 0.001$, to a very high and rather unrealistic vaccine unroll, $\alpha = 0.05$, so that we can have insights on the system behavior under extreme but meaningful

vaccination bounds.

For $\alpha = 0.001$, we can appreciate how $\Delta n_A(\infty) \approx 0$ for the most part of $(\theta, n_A(0))$ -space (Figure 2a1), this meaning that α is very low to trigger a substantial vaccination coverage (Figure 2b1), remaining very close to null, and thus system's prevalence tends to be rather homogeneous and maximal across the parameter space (Figure 2c1). The poor vaccination coverage cannot bring down the system's endemic phase in any point of the parameter space. Notably increasing α to $\alpha = 0.005$, a different picture starts to emerge. The region where $\Delta n_A(\infty) \neq 0$ extends toward higher $n_A(0)$ and θ (Figure




FIG. 2. **System's behavior under homogeneous thresholds for the US multilayer network.** Results for normalized activation change $\Delta n_A(\infty)$ (Column **a**), vaccination coverage $v(\infty)$ (Column **b**), and prevalence $r(\infty)$ (Column **c**). In every panel, outcomes are explored in $(\theta, n_A(0))$ space. Every row from 1 to 4 shows the solution for a different vaccination scenario as defined by the vaccination rate α . Every point in the diagrams amounts to 25 network realizations times 50 dynamical realizations. White dashed line contours mark an approximate (due to finite size effects) separation between $\Delta n_A(\infty) = 0$ or $r(\infty) = 0$ and non-null values of the respective observables.

2a2). These behavioral changes lead to mass vaccine uptake when initial vaccine support $n_A(0)$ is high enough, independently of the activation threshold θ (Figure 2b2). Consequently, the disease’s prevalence decreases in that same region (Figure 2c2). Doubling α to reach $\alpha = 0.01$, leads to a qualitatively different picture when looking at the system’s prevalence (Figure 2c3). Vaccination coverage $v(\infty)$ is high enough in an important part of the system that yields a disease-free phase for medium-high $n_A(0)$ and also for lower values of this initial support as long as θ is kept very small. Finally, in the very high effort scenario, $\alpha = 0.05$, the trend continues, as the disease-free phase extends to lower values of $n_A(0)$ and higher values of θ .

Overall, for a fixed value of θ , moving from $n_A(0) = 0$ to $n_A(0) = 1$, means that vaccine hesitancy loses ground, VC increases and disease prevalence tends to zero. In high enough α scenarios, for low θ , the transitions occur in a more abrupt way, whereas for medium to high θ , the transitions are smoother. Varying θ with fixed $n_A(0)$ does not have a notable effect except where a phase separation exists. If the initial vaccine acceptance $n_A(0)$ is high enough, even $\theta \rightarrow 1$ has no effect on vaccination and prevalence. A giant component of vaccinated agents can emerge fast enough to avoid a sizeable outbreak. On the other hand, as $n_A(0)$ decreases (for a fixed α), critical values of θ appear that, if surpassed, can bring the system from a disease-free phase to an endemic phase. This critical threshold θ , however, can be pushed toward higher values if the vaccination campaign proceeds at faster rates.

Importantly, though, even in the most intensive vaccination scenario explored, the epidemiological impact could be sizeable if the initial support for vaccination is rather mild, and the threshold to be convinced is not very high. This is clearly a hindrance product of social contagion being regarded as a complex contagion process rather than a simple one, requiring greater support to diffuse a pro-active attitude towards vaccine uptake.

Complementarily, in the appendix B, we offer curves of the explored observables for varying θ with sections of constant $n_A(0)$, where the abrupt nature of the changes in $\Delta n_A(\infty)$, $v(\infty)$, and $r(\infty)$ is better appreciated.

B. Heterogeneous thresholds from surveys

In the second stage of this work, we inform our model with activation thresholds derived from real data collected from the surveys carried out by Lazer et al. [45], at the beginning of the COVID-19 vaccination campaign in the US. These surveys offer, among other features, insights into vaccination acceptance and hesitancy, segmented by age at the national level, and also account for the fraction of people within a certain attitude towards vaccine uptake at the state level.

The attitudes are categorical and include: “already vaccinated”, inclined to be vaccinated “as soon as possi-

ble”, “after at least some people I know”, “after most people I know”, and those who “would not get the COVID-19 vaccine”. It is from this threshold-based perspective and its heterogeneous distribution—varying according to others’ vaccination statuses—that we draw inspiration for our coupled epidemic-vaccination-opinion dynamics model.

Now, since the control parameters θ and $n_A(0)$ previously used are now fixed for every state, the only free parameter is the vaccination rate α , which we use to extensively explore the outcomes of the coupled dynamics. Taking full advantage of the vaccination attitude data for every state in the US, simulations are conducted for every US state. It is important to note that the assignment of different vaccination attitudes and their corresponding thresholds to agents in the network is done randomly.

1. System’s behavior under different vaccination scenarios

Similarly as before, we examine how $\Delta n_A(\infty)$, $v(\infty)$, and $r(\infty)$ behave across varying α values in Figure 3, panels a, b, and c, respectively.

We observe that $\Delta n_A(\infty)$ across states remains between 0.2 and 0.6 for the vaccination scenarios with the lowest effort. This separation diminishes as α increases, aligning within the range of 0.9 to 1.0 for scenarios with the highest vaccination rate α . Notably, no state exhibits $\Delta n_A(\infty) = 0$ in any scenario, implying, that all of them can trigger behavioral cascades that increase the number of convinced individuals and thus there is substantial potential for vaccination coverage to grow. Regarding vaccination coverage, we appreciate how in low α regimes, VC is primarily between 0.1 and 0.2, with the exception of Alaska (AK). For intermediate α values, representing more realistic vaccine uptake rates, the disparity in VC across states reduces but widens again as $\alpha \rightarrow 1$. Due to inherent differences in each state’s contact network but, more importantly, their θ distribution, the curves show different rates of growth and intersect at various vaccination scenarios. This phenomenon reflects the diversity in vaccine attitudes across states. As a specific example, the case of Massachusetts (MA) is especially interesting. In the lowest α regimes, MA occupies an average rank among states, with 13% of its population having already received the vaccine. However, with increasing α , the vaccination-opinion feedback loop is further stimulated. This facilitates the progressive vaccination of individuals in subsequent categories, specifically “soon” at 47%, and “someone” at 18%. Ultimately, this leads MA to emerge as one of the states with the highest vaccination coverage. The disease’s prevalence, consequently, is the highest for the lowest vaccination rates, spanning approximately from 0.25 to 0.4. These values rapidly decrease with increasing vaccine unroll, reaching a disease-free state beyond $\alpha \approx 0.05$. Alaska is correspondingly an exception here too; its prevalence is an outlier for the lowest α scenarios ($r(\infty) < 0.1$), and sizeable outbreaks

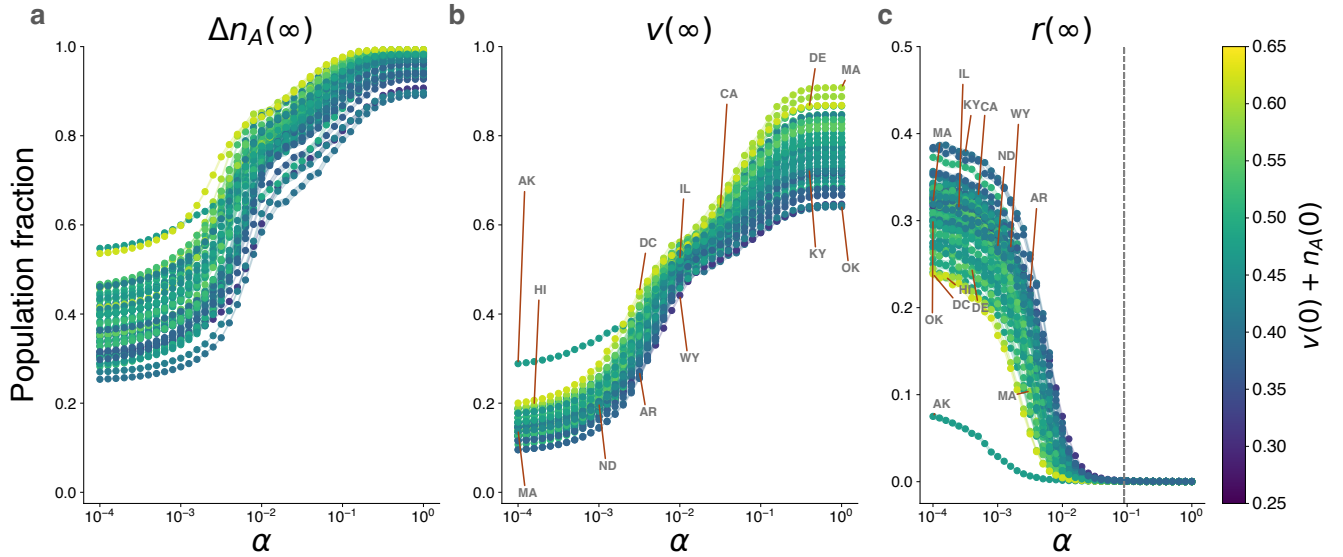


FIG. 3. **Epidemiological impact across US states for varying vaccination rates α .** The vertical dashed line represents an approximate separation between the endemic phase and the disease-free phase. Panel a: $\Delta n_A(\infty)$ vs. α . Panel b: $v(\infty)$ vs. α . Panel c: $r(\infty)$ vs. α . The color bar represents the initial fraction of support for the vaccine: “already vaccinated” $v(0)$ plus “as soon as possible” $n_A(0)$. Some state labels are shown for reference.

are negligible for $\alpha > 0.005$. With an initial “already vaccinated” fraction of $v(0) = 0.29$, Alaska accounts for the highest initial immunized population.

All the curves across the panels incorporate additional information represented by a color map for the sum of initial conditions for vaccinated and pro-vaccine individuals, $v(0) + n_A(0)$. Overall, there is a positive correlation between this initial condition and VC, and a negative one towards the disease’s prevalence.

2. Relevance of heterogeneous thresholds

To explore the effect of heterogeneity in vaccination attitudes, let us examine more closely the relevance of each vaccination attitude, and thus the threshold θ , in determining the final epidemic size $r(\infty)$.

To analyze this, we perform a pairwise correlation analysis between each state’s prevalence $r(\infty)$ and the fraction of the state’s population in different vaccination attitudes across the extensive range of vaccination scenarios examined earlier in this section (Figure 3), as determined by the vaccination rate α . Figure 4 shows Pearson’s correlation coefficient ρ for each vaccination category as a function of α . The main continuous lines represent the ρ values, while the error bars denote the 95% confidence intervals. Although the errors are relatively large due to the small number of data points (equal to the number of states, plus the District of Columbia, and the US level), the overall trends remain clear.

For the low to middle range of vaccination scenarios considered, $\alpha \in [10^{-4}, 6.5 \times 10^{-3}]$, the factor (vaccina-

tion category) that achieves greater explanatory power in terms of the Pearson coefficient is the ‘already’ category. That is, the initial condition $v(0)$ for the vaccinated population. In terms of our framework, they populate the initial condition $n_A(0)$ (together with those in the ‘soon’ category). Strictly speaking, though, this is not a behavioral attitude towards the vaccine that plays a role in the implemented opinion dynamics, but one of the categories in which individuals were classified when responding to the surveys in [45]. However, a varying value of $v(0)$ across states encapsulates heterogeneous predispositions of individuals in those regions as well as different vaccination efforts deployed in each state. As vaccination rates further increase α , a crossover appears between categories ‘already’ and ‘soon’, taking the ‘already’ category the dominant role in terms of explanatory power for a higher range of α values explored. When vaccination rates are high enough, states with a higher proportion in this vaccination category will see their vaccination coverage greatly increased due to this behavioral category for which vaccination is just a spontaneous process, since they are initially convinced independently of their social peers. People in the ‘someone’ category, on the other hand, would react to the vaccination process as if it were a simple contagion process: they just need an individual to change their attitude from hesitant to convinced and take the vaccine when available. For the most part, the contribution to explanatory power is very poor, especially in the lower α range. This is understandable for two reasons: i) according to the data (see Figure A1 in Appendix A), this is the lowest populated vaccination category across the majority of states, and ii) these




FIG. 4. **Pearson's coefficient (Prevalence vs. Population fraction) as a function of vaccination rates.** The coefficient is measured using a pairwise linear regression of each US state's average final prevalence versus the fraction of the population in the corresponding vaccination attitude category according to surveys. This measurement is performed for the considered range of vaccination scenarios as given by the vaccination rate α . Shaded areas cover the 95% C.I.

individuals require only one vaccinated neighbor to be convinced to vaccinate, thus they do not constitute a concerning bottleneck to trigger behavioral cascades.

The next vaccination attitude in terms of increasing hesitancy is 'most'. Individuals in this category are those who genuinely make their decision based on a complex contagion process for which $\theta > 0.5$. As we saw in the first part of this work when studying the system's behavior under a homogeneous activation threshold θ , a value of $\theta > 0.5$ can have an important epidemiological impact in the population unless the initial support for the vaccine $n_A(0)$ is sufficiently high, also conditioned to the vaccination effort α . According to the results in Figure 4, the impact of 'most' stays fairly constant, with the Pearson coefficient ρ around 0.5 for the whole range of α . The explanatory power is high, but other factors, either 'already' for low to middle α , or either 'soon' or 'never' for higher α show a greater relevance in these terms. What mitigates a potential negative impact of the 'most' category is the fact, as it happened with 'someone', that people in this category, although with a higher representation, are far from a critical fraction of the population

in their respective states. Finally, the 'never' category, the group of people who show the highest hesitancy or anti-vaccine sentiment, acquires the greatest relevance in the middle to the higher end of the explored range of α . With a high enough vaccine rollout, the states that will experience higher outbreaks will be those with a higher fraction of individuals who would not take the vaccine, logically blockading an almost full percolation of the vaccinated status across the network.

Indeed, we can appreciate by looking at the data that in several states the combined fractions of the 'most' and 'never' attitudes amount to less than 0.5 of the population. This indicates that in scenarios with rapid vaccination, individuals with a 'most' attitude are likely to find most of their neighborhood already vaccinated, prompting their shift to a pro-vaccine stance. Second, in most states, the fraction of the population already vaccinated or willing to be vaccinated as soon as possible, $v(0) + n_A(0)$, is more than 0.4 and often exceeds 0.5. Including the 'someone' category which, as we stated, requires just one vaccinated neighbor per individual, this vaccinated/pro-vaccine segment increases to over 0.5 in

most cases. This suggests that, from moderate α scenarios onwards, vaccination-opinion cascades can readily grow, leading to widespread vaccination coverage and significantly mitigated outbreaks.

IV. DISCUSSION

In this work, we have explored the interplay of contagion processes that differ in their transmission mechanisms: a classical pairwise SIR epidemic spread; and a vaccination model driven by a social complex contagion process based on the Watts-Granovetter threshold model. Far from an arbitrary theoretical problem, the model proposed is grounded in US surveys on vaccine hesitancy, which showed complex contagion features in the attitudes of individuals towards vaccine uptake [45].

We have characterized the system’s behavior by the relative size of the activation cascade $\Delta n_A(\infty)$, the vaccination coverage $v(\infty)$, and the disease’s prevalence $r(\infty)$ at the absorbing state in the control parameter space of $n_A(0)$, initial fraction of people convinced to take the vaccine, and θ , the activation threshold from hesitancy to convinced. First, under the assumption of homogeneous activation threshold θ across the population, within a realistic range of vaccination values, the emergence of a disease-free phase is possible for very low θ and medium $n_A(0)$. However, even under highly unrealistic vaccine rollout, sizeable outbreaks persist if the initial support $n_A(0)$ is not high enough and the activation threshold is moderate. In the second part, we have relaxed the homogeneity assumption by endorsing the population with heterogeneous thresholds based on real vaccine hesitancy surveys. We observed mixed behavior depending on the vaccination scenario. Generally, and as expected, individuals who are already vaccinated play a key role in reducing the epidemic’s final size. Conversely, those who significantly delay vaccine uptake or directly refuse to vaccinate have a negative impact on the disease’s prevalence. Fortunately, given that initial support for vaccine uptake is sufficiently high across states, the influence of anti-vaccine zealots is mitigated.

The results can be further contextualized by considering a few limiting regimes that help clarify the mechanisms at play. When the vaccination rate α is very low, individuals who become convinced to vaccinate do so too slowly to alter the epidemic trajectory. In this case, the system behaves similarly to a baseline SIR model, and the opinion dynamics have little epidemiological impact. In contrast, when α approaches its maximum value ($\alpha = 1$ in our discrete-time formulation), vaccination occurs almost immediately upon willingness, and the system becomes governed by the structure of the opinion dynamics—namely, the activation threshold θ and the initial support $n_A(0)$. A particularly interesting regime emerges when $n_A(0) \approx \theta$, as this defines the average level of vaccinated neighbors an agent would expect to observe at early times, assuming opinions are homoge-

neously distributed. In such a setting, behavioral cascades can be rapidly triggered if vaccination is fast, since most agents are just at the brink of activation. However, this potential tipping point is only effective when rollout is sufficiently fast—when α is low, this frontier does not mark a meaningful behavioral transition. Moreover, the observed transition is not sharply aligned with the $n_A(0) = \theta$ diagonal, likely due to network topology and higher-order structural features such as clustering. These subtleties highlight the interplay between social reinforcement, vaccination logistics, and contact structure in shaping vaccine adoption and epidemic outcomes.

Our model assumes that once an individual adopts a pro-vaccine stance, this decision is irreversible. This choice is consistent with the original formulation of the Watts-Granovetter threshold model, and aligns with our goal of integrating social reinforcement dynamics with epidemic spread in a focused, empirically grounded framework. Moreover, given the SIR nature of the epidemic process, such unidirectional opinion change has limited implications: once infected or recovered, individuals no longer participate in the vaccination process. Introducing opinion decay (e.g., $\mathcal{A} \rightarrow \mathcal{H}$) for susceptible individuals would likely slow the overall vaccination cascade and could worsen epidemic outcomes. Nonetheless, reversible opinion dynamics may become relevant in contexts with waning immunity (SIRS-like processes), slow vaccination rollout, or fluctuating perceived risk due to local incidence or media influence. Exploring such feedbacks—especially when conditioned on dynamic signals—would be a valuable extension of this framework.

One criticism that could be raised about our approach is the use of an age-multilayer contact structure while the dynamical processes running on it, as defined here, do not explicitly depend on age, nor do we analyze outcomes by age group. Our choice was to focus on global results in order to emphasize the interplay between epidemic dynamics and threshold-based opinion formation, rather than provide a more detailed epidemiological account. We also observe that the main conclusions hold for simpler, data-agnostic contact structures, such as homogeneous or heterogeneous single-layer networks (see Appendices D and E).

This modeling choice is closely related to another simplifying assumption: the random assignment of pro-vaccine stances across the population. However, vaccine hesitancy is known to vary by age, gender, or socioeconomic status [10, 54], and structured heterogeneities in these attributes may affect both vaccine coverage and epidemic outcomes, especially in populations with markedly different demographic profiles. Due to the lack of granular data on vaccination attitudes disaggregated by state and age, we have not pursued this more detailed stratification here, though it remains an important avenue for future research.

Moreover, incorporating community structure and homophily in the distribution of vaccine attitudes is critical to understanding the reemergence and successful trans-

mission of infectious diseases [55]. Behavioral clustering and assortative mixing in opinion networks—including the emergence of tightly knit anti-vaccine communities—are known to significantly influence both social reinforcement and epidemic outcomes [6, 28, 56, 57]. Our current model does not include these higher-order structural features, such as social clustering, modularity, or opinion assortativity, all of which can amplify or dampen behavioral cascades. Incorporating such features would provide a richer framework for exploring system resilience, vulnerability to misinformation, and the effectiveness of targeted interventions—particularly within the broader context of interacting simple biological and complex social contagion dynamics.

Finally, although the threshold model proposed is grounded in responses from real surveys, human decision-making is multifaceted, and additional mechanisms undoubtedly influence the choice to vaccinate. These mechanisms might include rational action, as is typically considered in vaccination game approaches [16, 20], where individuals weigh the costs of different options. Additionally, the influence of mass media and social networks is nowadays critically significant, where figures of authority or online communities may exert more influence than peers or acquaintances [58–61]. Moreover, these platforms can also facilitate the rapid spread of misinformation and fake news, further complicating the landscape of decision-making. Furthermore, communities evolve and adapt, often leading to fragmentation between opposing viewpoints and the creation of echo chambers [62–64]. These factors are increasingly pivotal in shaping attitudes toward vital issues and warrant consideration in future studies.

In summary, interesting potential extensions could be added to advance our understanding of the interplay between epidemics and social contagion processes related to vaccine uptake, but as it is, our model stands as an example of naturally competing simple and complex contagions processes, this last inspired by real-life observation of people’s social behavior from vaccine hesitancy surveys. This is valuable in itself given the scarcity of real examples of contagion processes where transmission is genuinely characterized by group interactions.

V. METHODS

A. Age-layered network

A population of N individuals is considered, structured in a multilayer network based on individuals’ age a . An individual i of age a belongs to layer- a , and their contacts k_i can be distributed both within this layer and across other layers, following realistic age-mixing patterns.

This contact structure is derived from contact matrices built from data-driven synthetic populations developed by Mistry et al. [46]. The contact matrices for different countries and regions—in particular, those used in this

work for the United States, District of Columbia, and the 50 states—are publicly available at [65]. These matrices encode detailed contact patterns, resolving up to 85 age groups.

It should be noted that these matrices were originally constructed using demographic data from 2005. To reflect a more current picture of American society, we use 2019 population census data [66] and update the contact matrices following the procedure in [67].

Let M be the original contact matrix, where M_{ab} denotes the average number of contacts between age group a and age group b during a fixed time period. The updated matrix M' is computed as:

$$M'_{ab} = M_{ab} \frac{NN'_b}{N_bN'}, \quad (7)$$

where N and N' are the total original and updated population sizes, and N_b and N'_b are the original and updated populations of age group b , respectively.

To generate the multilayer network \mathcal{M} from the contact matrix M (hereafter dropping the prime for clarity), we follow the procedure developed in [47]. Briefly:

1. For every node (individual) in a given layer a , stubs (half-links) are sampled from a negative binomial distribution, as observed in [47], with average degree in layer a given by:

$$\langle k \rangle_a = \sum_b M_{ab}.$$

2. The inter-layer probability matrix is computed as:

$$p_{ab} = \frac{M_{ab}}{\sum_b M_{ab}},$$

representing the probability that a stub from an individual in layer a connects to another in layer b .

3. With the degree k_i of each individual and the probabilities p_{ab} , a configuration model is implemented. Stubs are matched by uniformly sampling a source node i , selecting a target layer b according to p_{ab} , and sampling a target node j within that layer with probability proportional to their remaining stubs, i.e., $k_j / \sum_{\ell \in b} k_\ell$.

The result is a multilayer network \mathcal{M} where individuals are connected following realistic age-based contact patterns. However, it is important to note that this multilayer configuration model does not explicitly incorporate higher-order structural features such as clustering, or community structure.

As the network generation process is stochastic, final results from the dynamical simulations are averaged over an ensemble of network realizations generated under the same parameters.

B. Epidemiological model

The epidemiological model used is the standard and well-known susceptible-infected-removed SIR model running on a population of individuals structured on the age-based multilayer network as explained previously.

We consider discrete-time dynamics with unitary time step $\Delta t = 1$ day. Then, at every time step, susceptible individuals can experience a simple pairwise contagion interaction and become infected with probability $\beta\Delta t$:

$$\mathcal{S} + \mathcal{I} \xrightarrow{\beta} \mathcal{I} + \mathcal{I}. \quad (8)$$

Since every individual can interact with any other of their first neighbors in the network, the total probability of becoming infected is $P(\mathcal{S} \rightarrow \mathcal{I}) = 1 - (1 - \beta\Delta t)^{I_i(t)}$, where $I_i(t)$ denotes the number of infected neighbors of individual i at time t , and β is the disease's transmission rate. The quantity $I_i(t)$ is bounded by the node's degree k_i .

Lastly, recovery or removal from the infection dynamics is, as usual, a spontaneous transition modeled as a Poisson process:

$$\mathcal{I} \xrightarrow{\mu} \mathcal{R}. \quad (9)$$

Then, for every infected individual, the probability of decaying to the recovered/removed state is $P(\mathcal{I} \rightarrow \mathcal{R}) = \mu\Delta t$ at every time step.

C. Opinion dynamics model

The SIR process runs simultaneously with two other intimately related dynamical processes: the vaccination campaign and the opinion dynamics on vaccination. Vaccination by itself is modeled as a spontaneous process proceeding continuously as long as there are eligible individuals to be vaccinated and without any preferential target. The catch is that in order to be eligible for vaccine uptake, the individual must be in a convinced or active state. Here is where the opinion formation process enters.

To model the opinion dynamics we consider the Watts-Granovetter threshold model of social contagion [29, 31]. In this model, individuals can be classified into two states: inactive and active. Within this context, we will refer to the inactive ones as hesitant, being them denoted by \mathcal{H} , and we will retain the active nomenclature to refer to the pro-active ones regarding vaccine uptake, denoted by \mathcal{A} . As in the original model, the state \mathcal{A} is irreversible, and therefore the dynamics are ruled by the transition:

$$\mathcal{H} \rightarrow \mathcal{A}. \quad (10)$$

Mapping the status \mathcal{H} to 0 and \mathcal{A} to 1, and designating as o_i the individual's i opinion, the updating rule in the

original formulation of the model proceeds in the following way:

$$o_i(t+1) = \begin{cases} 1 & \text{if } \sum_{j \in \Omega_i} \frac{o_j(t)}{k_i} \geq \theta_i, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Here, the summation extends to Ω_i , which is the neighborhood of individual i , o_j is the status of the neighbor j , and k_i is the degree or number of neighbors of i ; finally, θ is the individual's activation threshold ($0 \leq \theta_i \leq 1$). Due to its lack of symmetry, if the focal individual status is already at 1, nothing happens. Thus, as aforementioned, once the state 1 is adopted, it is conserved until the end of the dynamical process.

For the present work, however, we do not follow the rule described in equation 11 but a variation. Rather than adopting a proactive attitude towards vaccination based on the vaccine views of peers, people will make the decision on vaccine uptake based on their neighbors' vaccination status. Therefore, the updating proceeds as follows:

$$o_i(t+1) = \begin{cases} \mathcal{A} & \text{if } \frac{V_{\Omega_i}(t)}{k_i} \geq \theta_i, \\ \mathcal{H} & \text{otherwise.} \end{cases} \quad (12)$$

Where $V_{\Omega_i}(t)$ is the total number of vaccinated individuals within i 's neighborhood. This change directly connects the success of the ongoing vaccination campaign to the opinion on vaccination, setting a co-evolving feedback loop between vaccination and opinion dynamics.

Finally, every convinced or active susceptible (\mathcal{S}, \mathcal{A}) will be vaccinated

$$\mathcal{S} \xrightarrow{\alpha} \mathcal{V}, \quad (13)$$

with probability $P(\mathcal{S} \rightarrow \mathcal{V}) = \alpha\Delta t$.

While we can find either infected or removed individuals with a pro-active vaccination stance because they were convinced just prior to following the infection route in our model, they will not be eligible for vaccination anymore, and thus will not pose any influence on the vaccination campaign.

D. Vaccine hesitancy data

In order to adapt the categorical attitudes to our numerical approach, as defined in the surveys by Lazer et al. [45], we translate them into quantifiable activation thresholds, denoted by θ . Individuals in the "already vaccinated" category are exempted from further opinion dynamics, effectively setting their activation threshold at zero. The same is presumed for individuals categorized as "as soon as possible". We posit that these individuals are already convinced and merely awaiting their opportunity to be vaccinated; hence, we assign them a $\theta = 0$. For those who would get vaccinated "after at least someone I know", we interpret this as requiring one

known individual to be vaccinated, leading to a threshold of $\theta = 1/k_i$, where k_i is the individual's number of first neighbors. Those in the “after most people I know” category are assigned a simple majority criterion, corresponding to $\theta = 0.5$. It could be argued that this approach is somewhat lenient with the constraint, considering that the majority criterion might vary on an individual basis. Lastly, for the “would not get the COVID-19 vaccine” category, we assign a threshold of $\theta = 1^+$, signifying that their decision remains unchanged irrespective of the vaccination status of their entire neighborhood, indicative of anti-vaccine sentiment. Table I presents a summary of the proposed mappings.

survey category	θ
<i>already vaccinated</i>	0
<i>as soon as possible</i>	0
<i>after at least someone I know</i>	$1/k_i$
<i>after most people I know</i>	0.5
<i>would not get the [COVID-19] vaccine</i>	1^+

TABLE I. **Vaccination attitude categories and activation threshold attribution.** Mapping to inform the individuals' activation thresholds in our model based on the vaccination attitude categories as given by Lazer et al. [45].

The specific fractions in every vaccination attitude category for every state in the US are depicted in Figure A1 of the supplementary material A

E. Simulations

Our results are obtained through extensive discrete-time Monte Carlo simulations. For the sake of clarification, we describe the algorithm followed to implement the coupled dynamics described above:

1. A multilayer network with N nodes is generated according to the method briefly described above and developed in [47]. The individual with the highest degree is assumed to be the patient zero, which is the initial seed that will trigger the epidemic spreading process. An initial fraction $n_A(0)$ of pro-vaccine individuals are randomly assigned across the network.
2. At every time step t , every node i in the contact multilayer network is visited, and the threshold model updating rule is applied (Eq. 12). In practice, a list of hesitant/inactive and susceptible individuals is kept since they are the only type of agents that could change their opinion status. Opinion updating follows a parallel scheme.
3. During the same time step, the transitions related to the SIR+V model occur. Susceptible and proactive individuals who were not vaccinated undertake a Bernoulli trial with each infected neighbor. Effective updating of health statuses also takes place following a parallel scheme.

4. If the population of infected individuals drops to zero, the process terminates, otherwise, we move to the next time step $t \rightarrow t + 1$ (and back to item 2).

In all the experiments performed, each result for a given set of control parameters has been derived by averaging over 25 network realizations and 50 iterations of the dynamical process for each network.

F. Code and data availability

The code developed for the data curation process, simulations, analysis of results, and figure generation, as well as the curated data to feed the model, is hosted at <https://github.com/phononautomata/threshold>.

ACKNOWLEDGMENTS

A.A. and Y.M. were partially supported by the Government of Aragon, Spain, and ERDF “A way of making Europe” through grant E36-20R (FENOL), and by Ministerio de Ciencia e Innovación, Agencia Española de Investigación (MCIN/AEI/ 10.13039/501100011033) Grant No. PID2023-149409NB-I00. A.A. acknowledges support through the grant RYC2021-033226-I funded by MCIN/AEI/10.13039/501100011033 and the European Union ‘NextGenerationEU/PRTR’.

Appendix A: Fraction of every vaccine attitude per state (survey data)

In figure A1, we report the fraction of the population with each vaccination attitude according to the survey by Lazer et al.[45].

In Figure A1, we report the fraction of the population in each vaccination attitude category according to the survey by Lazer et al.[45]. This survey provides valuable insights into public sentiment regarding vaccination across various states in the United States. The data are broken down into five distinct attitudes toward vaccination:

- “Already vaccinated”: This category includes individuals who have already received the COVID-19 vaccine.
- “As soon as possible”: These individuals are ready to get vaccinated as soon as they are eligible, reflecting a proactive approach toward vaccination.
- “After at least some people I know”: These individuals are hesitant but open to getting vaccinated once they observe others around them, such as friends or family, receiving the vaccine.
- “After most people I know”: This group is even more hesitant and prefers to wait until a larger portion of their social circle gets vaccinated.
- “I would never take the [COVID-19] vaccine”: The most resistant group, comprising individuals who are strongly opposed to taking the COVID-19 vaccine under any circumstances.

Appendix B: Complementary results for homogeneous thresholds

Complementarily to the heatmap shown in Figure 2, we represent in Figure B1 the system’s behavior for the US multilayer network as sections of constant $n_A(0)$ for varying θ under chosen vaccination scenarios.

Now the first row (Panels a1-a4) represents the evolution of $\Delta n_A(\infty)$ for the proposed vaccination scenarios in the main text. As α increases, we observe how $\Delta n_A(\infty)$ increases too, being the section with higher $n_A(0)$ those that sustain higher $\Delta n_A(\infty)$ values for a larger range of increasing θ . Whereas the lowest vaccination scenario cannot trigger a substantial behavioral cascade, since all the curves quickly decay to 0 for low θ , higher α provokes successful behavioral cascades as confirmed by looking at the evolution of $v(\infty)$ curves in the second row (Panels b1-b4). The greatest change among the considered scenarios is observed when increasing effort from $\alpha = 0.001$ to $\alpha = 0.005$, with $v(\infty) \approx 0.1$ changing to $v(\infty) > 0.5$, for the extreme case of $n_A(0) = 1$. Even for sections with a high level of initial hesitancy (low $n_A(0)$), the change

in $v(\infty)$ is relevant. Further increases of α naturally lead to higher $v(\infty)$, but the gains are clearly sublinear. As for the disease’s prevalence (Panels c1 to c4), we observe, as reported in the main text, a very high (and narrow) impact for the lowest vaccination scenario, and as we increase α , we appreciate the relevance of a different $n_A(0)$, and the fact that a disease-free phase is reached when this initial support $n_A(0)$ is low enough (clearly the transition occurs between $\alpha = 0.005$ and $\alpha = 0.1$).

Finally, these sections allow us to clearly appreciate the abrupt nature of the changes in $\Delta n_A(\infty)$, $v(\infty)$ and $r(\infty)$ for varying θ . Interestingly, as α increases, changes are more abrupt, as the competition between the different dynamic processes involved intensifies.

Appendix C: Complementary results for surveys’ heterogeneous thresholds

Figure C1 depicts the specific scatter plots of prevalence versus population fraction in vaccination attitude categories for three selected vaccination scenarios ($\alpha = 0.001$, 0.005 , and 0.01).




FIG. A1. **US states' population fractions by attitude towards vaccination.** From the bottom to the top, for each state the bar segments represent the following attitudes towards vaccination: “already vaccinated”, “as soon as possible”, “after at least some people I know”, “after most people I know”, “I would never take the [COVID-19] vaccine”, as obtained from Lazer et al. [45].




FIG. B1. **Sections for the US multilayer network under homogeneous thresholds.** Panels 1 to 4 represent different vaccination efforts ($\alpha = 0.001, 0.005, 0.01, 0.05$ from left to right columns). Panels **a** depict $\Delta n_A(\infty)$, panels **b** represent $v(\infty)$, and **c**, $r(\infty)$ curves for constant $n_A(0)$ (value given in the color bar) and varying θ .

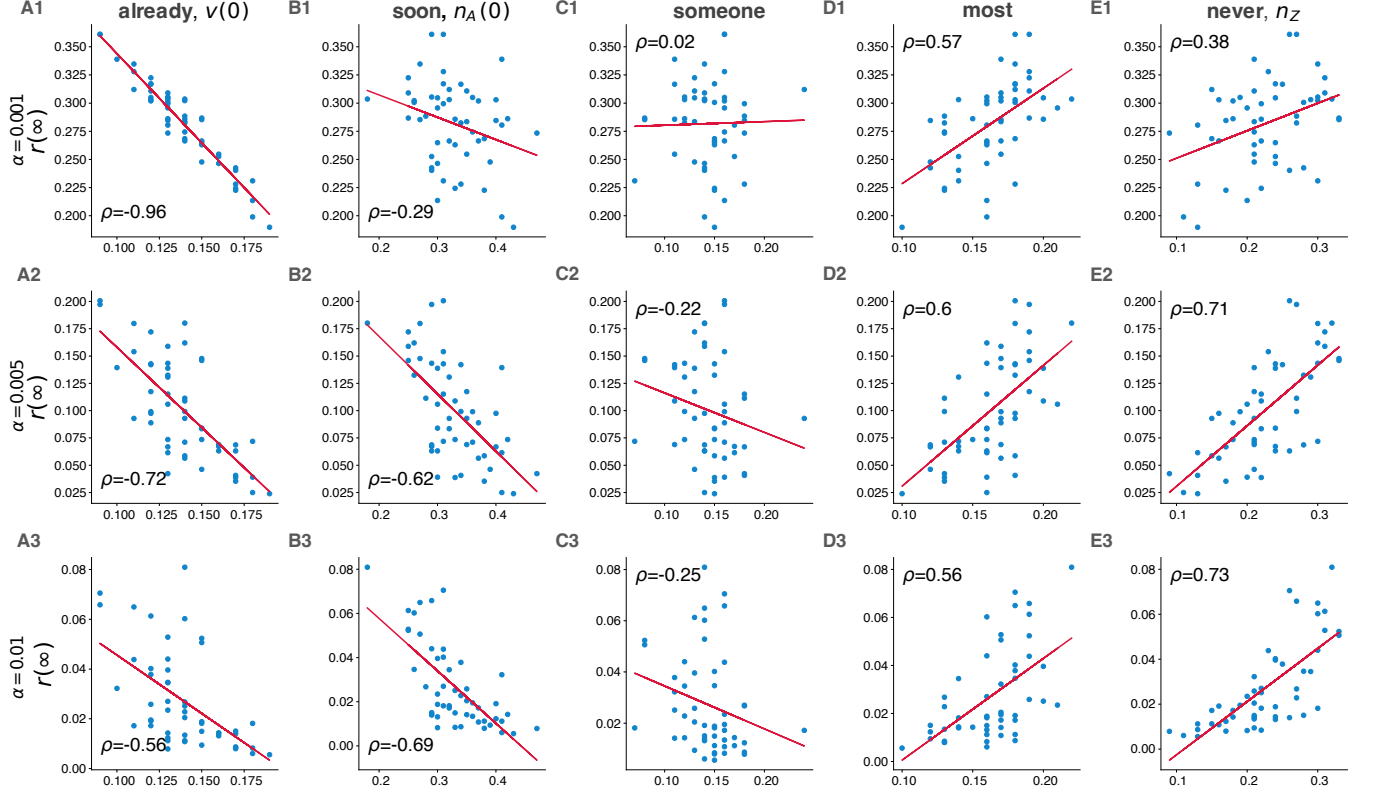


FIG. C1. **Pairwise scatter plots of prevalence and vaccine attitudes.** Panels with letters A, B, C, D, and E show, respectively, the scatter plots of prevalence (y-axis) against vaccine attitudes (x-axis): “already vaccinated”, “as soon as possible”, “after at least some people I know”, “after most people I know”. “I would never take the vaccine”. The state of Alaska has been removed as an outlier due to the markedly lower epidemiological impact due to the fact that its high initial condition $v(0)$.

Appendix D: Homogeneous thresholds on simple Erdős-Rényi and Barabási-Albert networks

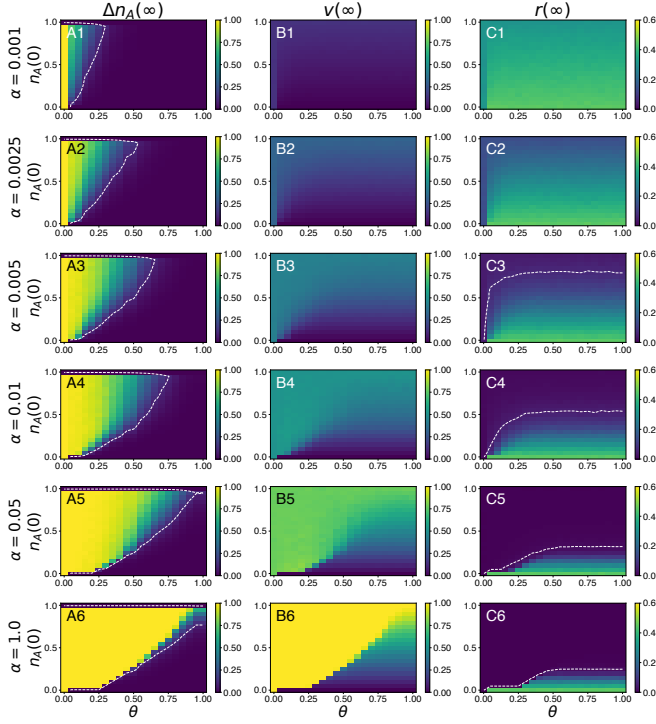


FIG. D1. **Homogeneous thresholds in Erdős-Rényi networks.** Results for normalized activation change $\Delta n_A(\infty)$, vaccination coverage $v(\infty)$, and prevalence $r(\infty)$. In every panel, outcomes are explored in $(\theta, n_A(0))$ space. Every row shows the solution for a different vaccination rate α . Every point in the diagrams amounts to 25 network realizations times 25 dynamical realizations. White dashed line contours mark an approximate (due to finite size effects) separation between $\Delta n_A(\infty) = 0$ or $r(\infty) = 0$ and non-null values of the respective observables.

As a complement to our main text results, we also show results of the same coupled dynamics under a homogeneous activation threshold θ in Erdős-Rényi and Barabási-Albert networks both with average degree $\langle k \rangle = 10$.

1. Erdős-Rényi networks

Although not presented here, the trivial scenario with $\alpha = 0$ triggers no cascading behavior, and subsequently, no vaccination at all (since $v(0) = 0$). Therefore the prevalence is always the maximum expected according to the given disease and network characteristics. In this case, for the stochastic simulations on the ER network, we obtain an average value $r_{ER}^{\max} \equiv r_{ER}(\infty) = 0.454$, with $[0.450, 0.460]$ as the 95% confidence interval (CI). Indeed, for the limit case $n_A(0) \rightarrow 0$ and $\theta > 0$, with $\theta > n_A(0)$, vaccination is trivially null for any α and

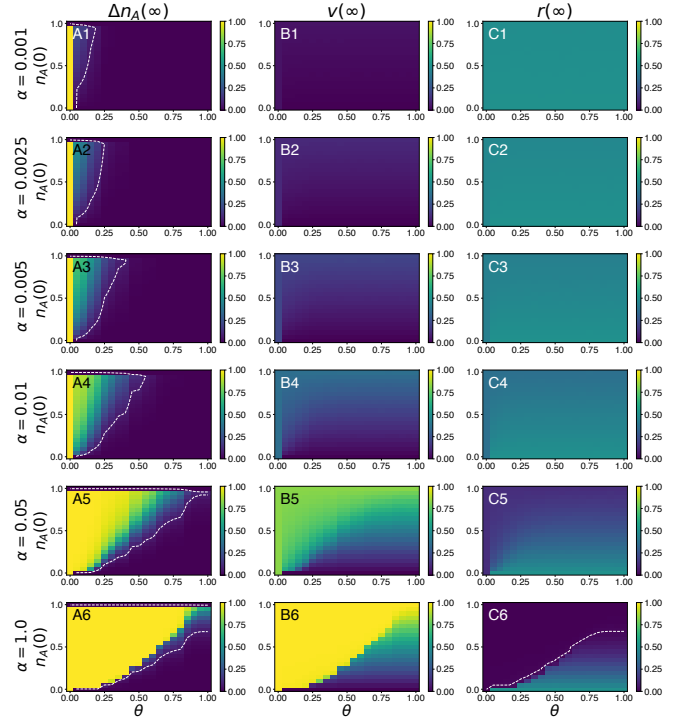


FIG. D2. **Homogeneous thresholds in Barabási-Albert networks.** Results for normalized activation change $\Delta n_A(\infty)$, vaccination coverage $v(\infty)$, and prevalence $r(\infty)$. In every panel, outcomes are explored in $(\theta, n_A(0))$ space. Every row shows the solution for a different vaccination rate α . Every point in the diagrams amounts to 25 network realizations times 25 dynamical realizations. White dashed line contours mark an approximate (due to finite size effects) separation between $\Delta n_A(\infty) = 0$ or $r(\infty) = 0$ and non-null values of the respective observables.

the system's prevalence approaches the case with $\alpha = 0$, $r_{ER}(\infty) \rightarrow r_{ER}^{\max}(\infty)$. As a further reference, given the values being used as epidemiological parameters, $\beta = 0.03$ for the transmission rate and $\mu = 0.2$ for the removal rate, and the social interaction term $\langle k \rangle = 10$, under a homogeneous mixing approximation we would expect $R_0 = \beta \langle k \rangle / \mu = 1.5$, and thus $r_{hom}(\infty) = 0.58$ as computed from the classical SIR solution given by the transcendental equation $r(\infty) = 1 - e^{-R_0 r(\infty)}$. Despite ER networks being regarded as homogeneous systems that can resemble more a standard well-mixing approach than, for instance, scale-free networks where the degree distribution has a power-law form, the relatively low $\langle k \rangle$ results in a significant contact saturation that leads to a reduced epidemiological impact[68].

For the lowest proposed vaccination rate scenario, $\alpha = 0.001$, we observe that for the most of the $(\theta, n_A(0))$ -space, there is no growing support for vaccine adoption, that is, $\Delta n_A(\infty) = 0$ (Figure D1 panel A1). We refer to the space where $\Delta n_A(\infty) = 0$ to the *impassive* region. Consequently, vaccination coverage is extremely low (Figure D1 panel B1), and sizeable outbreaks emerge,

ranging roughly between $r(\infty) \in (0.3, 0.46)$ (Figure D1 panel C1). The region of the parameter space where adherence to vaccine uptake is the maximum possible, $\theta = 0$ with $n_A(0) \in (0, 1)$, shows a slightly larger vaccination coverage and a lower prevalence. In any case, all the parameter space is characterized by an endemic solution ($r(\infty) \neq 0$).

As the vaccination rate α increases, there is a growing region of $(\theta, n_A(0))$, where hesitancy recedes and fades away, and consequently vaccine uptake emerges and the disease prevalence decreases. For $\alpha = 0.0025$, the maximum VC attained is at $v(\infty) = 0.324$ [0.318, 0.327] (Figure D1 panel B2), and the corresponding prevalence is reduced to $r(\infty) = 0.145$ [0.140, 0.149] (Figure D1 panel C2). The impact has been notably reduced in some regions of the parameter space, but the vaccination rates are still slow enough to avoid the emergence of epidemic outbreaks. Doubling α , $\alpha = 0.05$, clearly, the solution landscape changes. We attain a vaccination coverage around $v(\infty) = 0.425$ [0.417, 0.432] (Figure D1 panel B3), which provokes the beginning of an emergent disease-free region (Figure D1 panel C3). Still increasing α values translates into successful positive feedback between the vaccination adoption opinion dynamics and the vaccination campaign, which in turn increases the parameter space region where a disease-free solution reigns and, consequently, the endemic solution is bounded to domains with high adoption threshold θ and low to very low initial support $n_A(0)$ (Figure D1 panels C4, C5, and C6, progressively).

Overall, for a fixed value of θ , moving from $n_A(0) = 0$ to $n_A(0) = 1$, means that hesitancy loses ground, vaccination coverage increases, and disease prevalence tends to zero. In high enough α scenarios, for low θ , the transitions occur in a more abrupt way, whereas for medium to high θ , the transitions are smoother. Varying θ with fixed $n_A(0)$ does not have a notable effect except where a phase separation exists. If the initial vaccine acceptance $n_A(0)$ is high enough, even $\theta \rightarrow 1$ has no effect on vaccination and prevalence. A giant component of vaccinated individuals can emerge fast enough to avoid a sizeable outbreak. On the other hand, as $n_A(0)$ decreases (for a fixed α), critical values of θ appear that, if surpassed, can bring the system from a disease-free phase to an endemic phase. This critical threshold θ , however, can be pushed toward higher values if the vaccination campaign proceeds at faster rates.

The extreme and unrealistic case of $\alpha = 1$, depicts an abrupt transition when looking at the behavior of $\Delta n_A(\infty)$ in $(\theta, n_A(0))$ -space (Figure D1 panel A6). A rather marked boundary separates the region where the maximum size of cascading behavior occurs and the region where factoring in the initial support $n_A(0)$, there is no further change induced. This *impassive* region precludes the system from reaching total VC there. However, for the largest part of it (roughly when $n_A(0) > 0.25$), VC, as being propelled by a high vaccination rate, is large enough to drive the system to the free-disease

phase. Finally, it is noteworthy that the solution landscape for $r(\infty)$ is very similar across the vaccination scenarios with $\alpha = 0.05$ and $\alpha = 1$ (Figure D1 panels C5 and C6, respectively). This observation suggests the existence of diminishing returns to the vaccination process as determined by α , possibly constrained by the underlying network topology. Therefore, a moderate vaccination may be sufficient to efficiently mitigate the epidemic impact of the disease.

2. Barabási-Albert networks

For $\alpha = 0.001$, the situation is qualitatively similar to the coupled dynamics running on homogeneous networks. Vaccination rates are too low to trigger a cascading behavior and change individuals' hesitant opinions toward vaccine-proactive individuals. Virtually, VC is null across the $(\theta, n_A(0))$ (Figure D2 panel B1), and prevalence is maximal for the BA network under the current epidemiological parameters, being $r_{BA}(\infty) = 0.504$ [0.503, 0.505] (Figure D2 panel C1). However, at odds with the ER network case, continuing to increase α has hardly an effect. Adherence to vaccination advances very slowly (Figure D2 panels A1 to A4), and consequently VC is sluggish. Even at $\alpha = 0.01$, disease prevalence continues to be very high across all the $(\theta, n_A(0))$ space. Still, even at $\alpha = 0.05$, an important region of the parameter space suffers a high epidemic impact, and no disease-free solution emerges. Lastly, for $\alpha = 1$, the situation of $\Delta n_A(\infty)$ (Figure D2 panel A6) and $v(\infty)$ (panel B6) now resembles more closely the results for the ER network (Figure D1 panels A6 and B6, respectively). However, in the impassive region ($\Delta n_A(\infty) = 0$) where VC remains very low to null, the epidemic impact is larger than in homogeneous networks.

In BA networks, the competition between disease contagion and the opinion-vaccination dynamics is overwhelmingly won by the former. This outcome is expected, considering the well-documented fact that the topology of heterogeneous networks facilitates a rapid spread of disease [69]. In this coupled dynamical system under study, disease transmission is a simple contagion process, requiring only a single contact between an infected and a susceptible individual to propagate, whereas opinion dynamics—a complex contagion—demands a critical mass of individuals who are not only persuaded but also vaccinated, to initiate a cascading effect of behavioral change. In contrast to the dynamics on ER networks, the heavy-tailed degree distribution found in BA networks means that highly connected nodes play a significant role and necessitate a larger number of influenced neighbors to shift their stance toward a pro-vaccination viewpoint, thereby impeding or delaying widespread vaccination coverage.

Appendix E: Survey-based thresholds on ER and BA networks

1. System’s behavior under different vaccination scenarios

To finalize the analysis on ER and BA networks, we extend the model to the survey-based thresholds scenario as we did in the main text for the data-driven multilayer.

In Panels B1 and B2 of Figure E1, we focus on vaccination coverage. In low α regimes, VC is primarily between 0.1 and 0.2, with the exception of Alaska (AK). For intermediate α values, representing more realistic vaccine uptake rates, the disparity in VC across states reduces but widens again as $\alpha \rightarrow 1$. Due to inherent differences in each state’s $\langle k \rangle$ and, more importantly, their θ distribution, the curves intersect at various points. These varying gaps and curve crossings for different α values reflect the diversity in vaccine attitudes across states. As a specific example, the case of Massachusetts (MA) is interesting. In the lowest α regimes, MA occupies a median rank among states, with 13% of its population having already received the vaccine. However, with increasing α , the vaccination-opinion feedback loop is further stimulated. This facilitates the progressive vaccination of individuals in subsequent categories, specifically “soon” at 47%, and “someone” at 18%. Ultimately, this leads MA to emerge as one of the states with the highest vaccination coverage.

Regarding the final disease prevalence $r(\infty)$, Panels C1 and C2 display results for ER and BA networks, respectively. The most notable difference between the network types aligns with previous analyses. For ER networks (Figure E1 C1), in very low α regimes, prevalence across US states ranges from 0.1 to 0.3. These values decrease rapidly with increasing vaccination deployment, reaching a disease-free state beyond $\alpha \approx 0.05$. Alaska is correspondingly an exception here too; its prevalence is negligible across all vaccination rate scenarios, as represented by α . With an initial “already vaccinated” fraction of $v(0) = 0.29$, Alaska’s outcomes suggest that for very low α , the opinion-vaccination cascades do not initiate, resulting in $v(\infty) \approx v(0)$. This indicates that such an initial $v(0)$ is close to the herd immunity threshold p_c , preventing sizeable outbreaks. For a well-mixed population, akin to an ER network with a basic reproduction number $R_0 = 1.5$, the herd immunity threshold is $p_c = 1 - 1/R_0 \approx 0.33$. Alaska’s negligible impact is thus understandable in light of its initial condition $v(0)$ [70]. Conversely, in BA networks we saw the impossibility of attaining a disease-free solution for most α values. Now, even Alaska experiences significant outbreaks for most α values. In the case of Massachusetts, previously highlighted for its vaccination coverage (VC) progression with varying α in both ER and BA networks, a similar trend is evident in its prevalence data. It begins at a mid-rank position but becomes one of the states least impacted when α reaches high (yet realistic) values.

Overall, while the relationship is not strictly monotonic, it is evident that larger $v(0) + n_A(0)$ values correspond to higher VC and lower prevalence, and vice versa, underscoring the significance of initial vaccine support conditions in influencing the epidemiological impact.

2. Relevance of heterogeneous thresholds

We explicitly show scatter plots for the different vaccination attitude categories and each US state’s final epidemic size for selected vaccination scenarios (see Figure E2).

We observe a clear correlation as indicated by Pearson’s correlation coefficient for low and moderate α scenarios of $r(\infty)$ and $v(0)$ (“already vaccinated”) (refer to Figure E2 A1 to A6) and a relatively small or very poor correlation with respect to the “soon” category (Figure E2 B1 to B6), the only exception being $\alpha = 0.1$ for the BA network case (Panels A6 and B6). As expected, the higher $v(0)$ is, the lower $r(\infty)$. Particularly in the lowest vaccination rate scenarios, the correlation between $r(\infty)$ and $v(0)$ is very high (Panel A1 with $\rho = -0.97$) or even absolute (Panel A4 with $\rho = -1$), in contrast to $r(\infty)$ vs. $n_A(0)$. In these scenarios, where the vaccination campaign is irrelevant, it is evident that the emphasis is on the already vaccinated portion of the population rather than those pro-vaccine but not yet vaccinated. We observe that as α increases, the Pearson coefficient’s explanatory power for $v(0)$ diminishes, while it grows (in absolute value) for $n_A(0)$.

For the attitude “someone” (I will take the vaccine after someone I know), we observe mixed results. At the lowest α , it indeed correlates positively with prevalence, and as α increases, the sign of the correlation changes, and ρ grows in magnitude. However, correlations are generally weak, suggesting that these individuals are not key in shaping the spread.

The “most” attitude (I will take the vaccine after most people I know) shows moderate but distinct correlations. In states where this fraction of individuals is larger, larger outbreaks tend to occur. Across different α scenarios and networks tested, the correlation coefficient ρ varies within a small range. One reason these individuals are not more influential is that if a cascade of pro-vaccine individuals grows quickly enough, those with the “most” attitude will also likely become pro-vaccine, thus helping to reduce the final epidemic size. On the other hand, the presence of such individuals would be highly detrimental if it were to be accompanied by either high anti-vaccine sentiment or low pro-vaccine attitudes.

Finally, we consider the anti-vaccine attitude population fraction (“never”), i.e. the zealots, denoted by n_Z . A higher n_Z correlates with a higher $r(\infty)$, although their statistical correlation varies depending on the vaccination scenario. Similar to other attitudes, the fraction of zealots or anti-vaccine individuals increases its explanatory power under more intensive vaccination campaigns.

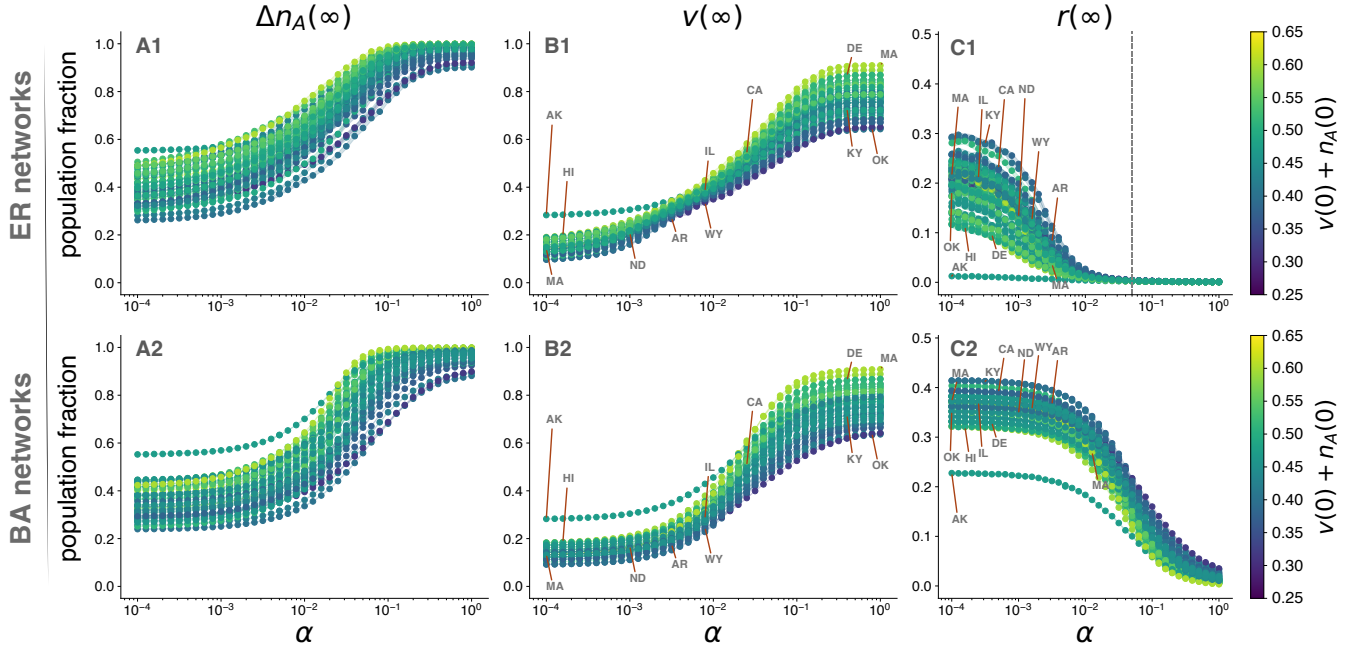


FIG. E1. **Epidemiological impact across US states for varying vaccination rates α .** Top panels: Results across US states on ER networks. The vertical dashed line represents an approximate separation between the endemic phase and the disease-free phase. Bottom panels: BA networks. Panels A1 and A2: $\Delta n_A(\infty)$ vs. α . Panels B1 and B2: $v(\infty)$ vs. α . Panels C1 and C2: $r(\infty)$ vs. α . The color bar represents the initial fraction of support for the vaccine: “already vaccinated” $v(0)$ plus “as soon as possible” $n_A(0)$. Some state labels are shown for reference.

Indeed, when vaccination becomes less relevant, it is logical to expect that individuals refusing vaccines under any circumstances will not have a significant impact. Their influence becomes particularly notable in BA networks for high to very high vaccination rates.

Although not directly comparable, assuming a uniform threshold setting in ER networks with initial vaccination support between 0.4 and 0.6, we see in Figure D1 (Panel C5) that we are within the disease-free phase of the diagram. Similarly, Figure E1 (Panel C1) shows that all states reach a disease-free solution at approximately $\alpha = 0.05$. This suggests that the initial conditions across states are favorable enough to effectively yield reduced prevalence, provided vaccination progresses swiftly. Conversely, for BA networks, consistently with the simplified homogeneous setting previously explored, even such robust initial support is insufficient to significantly lower the system’s prevalence.

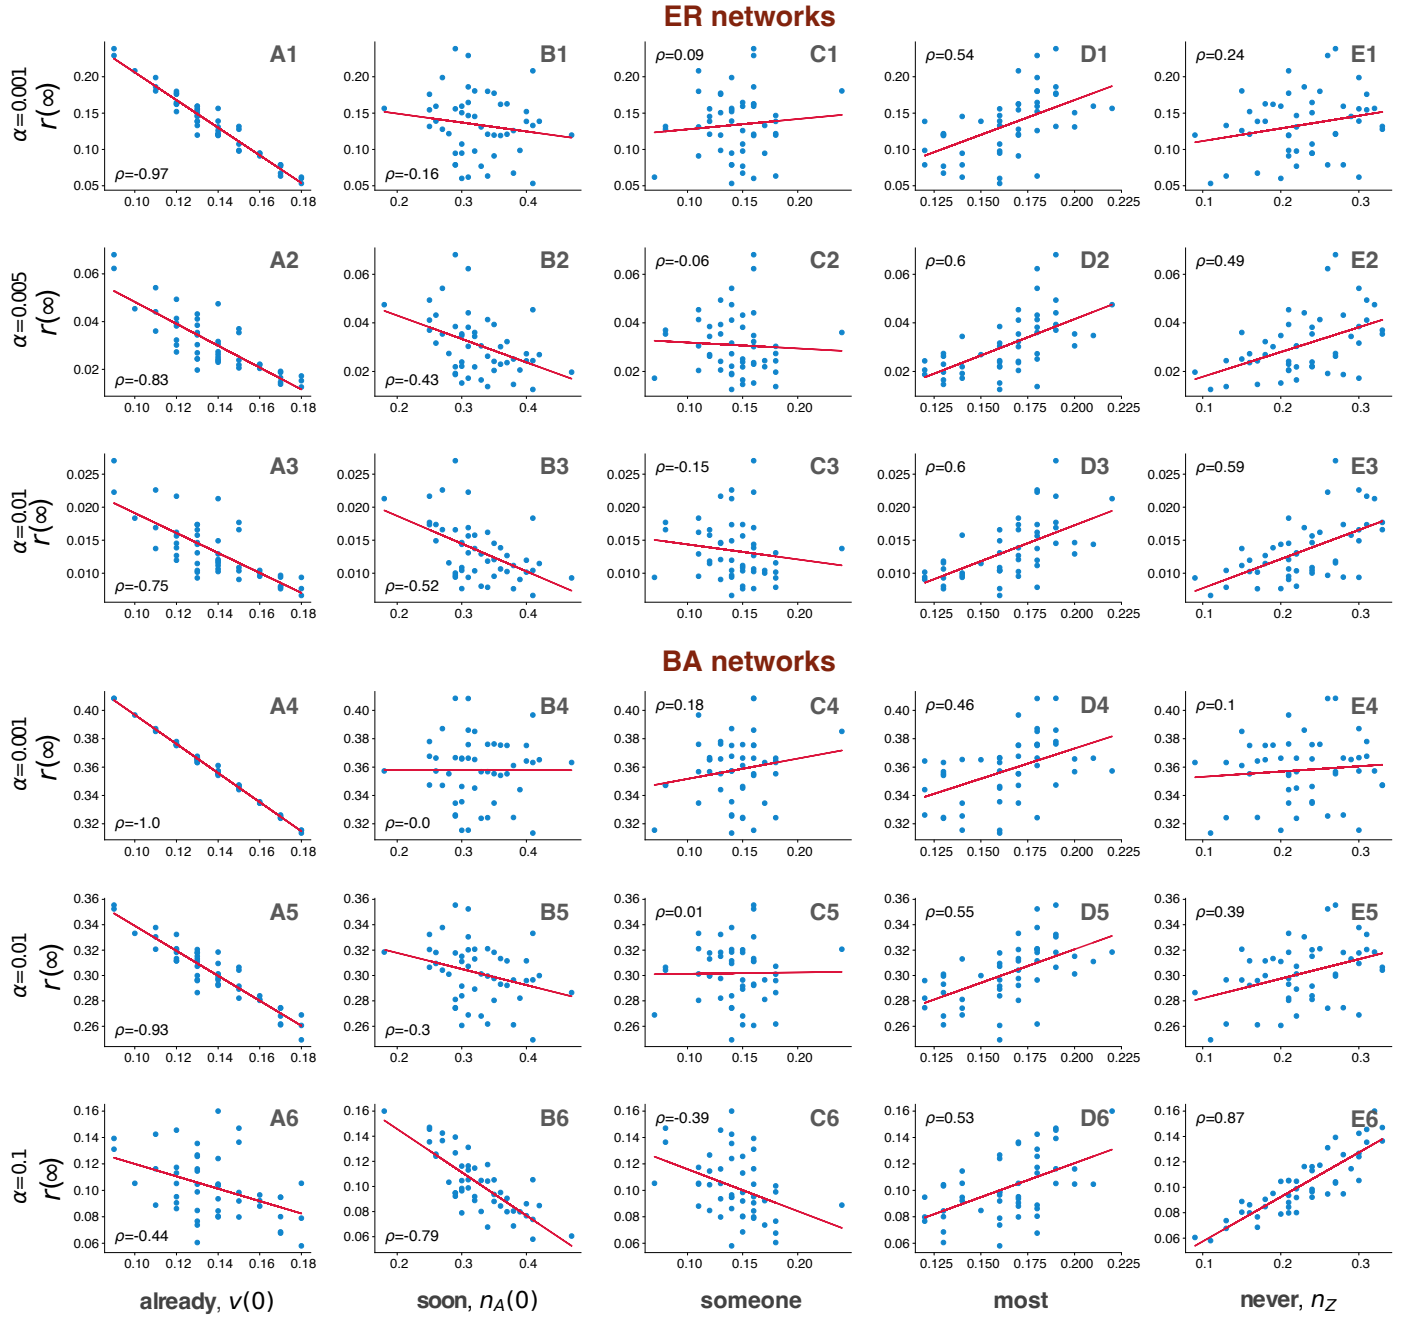


FIG. E2. **Pairwise scatter plots of prevalence and vaccine attitudes.** Panels with letters A, B, C, D, and E show, respectively, the scatter plots of prevalence (y-axis) against vaccine attitudes (x-axis): “already vaccinated”, “as soon as possible”, “after at least some people I know”, “after most people I know”. “I would never take the vaccine”.

- [1] N. Ferguson, Capturing human behaviour, *Nature* **446**, 733 (2007).
- [2] S. Funk, M. Salathé, and V. A. Jansen, Modelling the influence of human behaviour on the spread of infectious diseases: a review, *Journal of the Royal Society Interface* **7**, 1247 (2010).
- [3] N. Perra, D. Balcan, B. Gonçalves, and A. Vespignani, Towards a characterization of behavior-disease models, *PloS one* **6**, e23084 (2011).
- [4] R. Vardavas, R. Breban, and S. Blower, Can influenza epidemics be prevented by voluntary vaccination?, *PLoS computational biology* **3**, e85 (2007).
- [5] J. M. Epstein, J. Parker, D. Cummings, and R. A. Hammond, Coupled contagion dynamics of fear and disease: mathematical and computational explorations, *PloS one* **3**, e3955 (2008).
- [6] M. Salathé and S. Bonhoeffer, The effect of opinion clustering on disease outbreaks, *Journal of The Royal Society Interface* **5**, 1505 (2008).
- [7] L. Mao and Y. Yang, Coupling infectious diseases, human preventive behavior, and networks—a conceptual framework for epidemic modeling, *Social science & medicine* **74**, 167 (2012).
- [8] A. Rizzo, M. Frasca, and M. Porfiri, Effect of individual behavior on epidemic spreading in activity-driven networks, *Physical Review E* **90**, 042801 (2014).
- [9] E. Dubé, C. Laberge, M. Guay, P. Bramadat, R. Roy, and J. A. Bettinger, Vaccine hesitancy: an overview, *Human vaccines & immunotherapeutics* **9**, 1763 (2013).
- [10] G. Troiano and A. Nardi, Vaccine hesitancy in the era of covid-19, *Public health* **194**, 245 (2021).
- [11] D. Bernoulli and D. Chapelle, Essai d’une nouvelle analyse de la mortalité causée par la petite vérole, et des avantages de l’inoculation pour la prévenir (1760).
- [12] H. W. Hethcote, An immunization model for a heterogeneous population, *Theoretical population biology* **14**, 338 (1978).
- [13] X. Liu, Y. Takeuchi, and S. Iwami, Svir epidemic models with vaccination strategies, *Journal of Theoretical biology* **253**, 1 (2008).
- [14] Y.-H. Hsieh, Age groups and spread of influenza: implications for vaccination strategy, *BMC infectious diseases* **10**, 1 (2010).
- [15] Z. Wang, C. T. Bauch, S. Bhattacharyya, A. d’Onofrio, P. Manfredi, M. Perc, N. Perra, M. Salathé, and D. Zhao, Statistical physics of vaccination, *Physics Reports* **664**, 1 (2016).
- [16] Z. Wang, M. A. Andrews, Z.-X. Wu, L. Wang, and C. T. Bauch, Coupled disease–behavior dynamics on complex networks: A review, *Physics of life reviews* **15**, 1 (2015).
- [17] Q. Yin, Z. Wang, C. Xia, and C. T. Bauch, Impact of co-evolution of negative vaccine-related information, vaccination behavior and epidemic spreading in multilayer networks, *Communications in Nonlinear Science and Numerical Simulation* **109**, 106312 (2022).
- [18] S. L. Chang, M. Piraveenan, P. Pattison, and M. Prokopenko, Game theoretic modelling of infectious disease dynamics and intervention methods: a review, *Journal of Biological Dynamics* **14**, 57 (2020).
- [19] J. Tanimoto, *Sociophysics approach to epidemics*, Vol. 23 (Springer, 2021).
- [20] Y. Iwamura and J. Tanimoto, Realistic decision-making processes in a vaccination game, *Physica A: Statistical Mechanics and its Applications* **494**, 236 (2018).
- [21] C. Castellano, S. Fortunato, and V. Loreto, Statistical physics of social dynamics, *Reviews of modern physics* **81**, 591 (2009).
- [22] S. Ni, W. Weng, and H. Zhang, Modeling the effects of social impact on epidemic spreading in complex networks, *Physica A: Statistical Mechanics and its Applications* **390**, 4528 (2011).
- [23] L. G. Alvarez-Zuzek, C. E. La Rocca, J. R. Iglesias, and L. A. Braunstein, Epidemic spreading in multiplex networks influenced by opinion exchanges on vaccination, *PloS one* **12**, e0186492 (2017).
- [24] M. A. Pires and N. Crokidakis, Dynamics of epidemic spreading with vaccination: impact of social pressure and engagement, *Physica A: Statistical Mechanics and its Applications* **467**, 167 (2017).
- [25] M. A. Pires, A. L. Oestereich, and N. Crokidakis, Sudden transitions in coupled opinion and epidemic dynamics with vaccination, *Journal of Statistical Mechanics: Theory and Experiment* **2018**, 053407 (2018).
- [26] M. A. Pires, A. L. Oestereich, N. Crokidakis, and S. M. D. Queirós, Antivax movement and epidemic spreading in the era of social networks: nonmonotonic effects, bistability, and network segregation, *Physical Review E* **104**, 034302 (2021).
- [27] M. Fügenschuh and F. Fu, Overcoming vaccine hesitancy by multiplex social network targeting, in *International Conference on Complex Networks and Their Applications* (Springer, 2022) pp. 576–587.
- [28] A. Teslya, H. Nunner, V. Buskens, and M. E. Kretzschmar, The effect of competition between health opinions on epidemic dynamics, *PNAS nexus* **1**, pgac260 (2022).
- [29] M. Granovetter, Threshold models of collective behavior, *American journal of sociology* **83**, 1420 (1978).
- [30] D. J. Watts, A simple model of global cascades on random networks, *Proceedings of the National Academy of Sciences* **99**, 5766 (2002).
- [31] D. J. Watts and P. S. Dodds, Influentials, networks, and public opinion formation, *Journal of consumer research* **34**, 441 (2007).
- [32] D. Centola and M. Macy, Complex contagions and the weakness of long ties, *American journal of Sociology* **113**, 702 (2007).
- [33] B. Min and M. San Miguel, Competing contagion processes: Complex contagion triggered by simple contagion, *Scientific reports* **8**, 10422 (2018).
- [34] G. Cencetti, D. A. Contreras, M. Mancastroppa, and A. Barrat, Distinguishing simple and complex contagion processes on networks, *Physical Review Letters* **130**, 247401 (2023).
- [35] M. Salathé and S. Khandelwal, Assessing vaccination sentiments with online social media: implications for infectious disease dynamics and control, *PLoS computational biology* **7**, e1002199 (2011).
- [36] E. Campbell and M. Salathé, Complex social contagion makes networks more vulnerable to disease outbreaks, *Scientific reports* **3**, 1905 (2013).
- [37] F. Fu, N. A. Christakis, and J. H. Fowler, Dueling bio-

- logical and social contagions, *Scientific reports* **7**, 43634 (2017).
- [38] D. Centola, The complex contagion of doubt in the anti-vaccine movement, in *Meeting the challenge of vaccination hesitancy* (The Aspen Institute Washington DC, 2020) pp. 126–136.
- [39] L. Hébert-Dufresne, D. Mistry, and B. M. Althouse, Spread of infectious disease and social awareness as parasitic contagions on clustered networks, *Physical Review Research* **2**, 033306 (2020).
- [40] Z. Qiu, B. Espinoza, V. V. Vasconcelos, C. Chen, S. M. Constantino, S. A. Crabtree, L. Yang, A. Vullikanti, J. Chen, J. Weibull, *et al.*, Understanding the coevolution of mask wearing and epidemics: A network perspective, *Proceedings of the National Academy of Sciences* **119**, e2123355119 (2022).
- [41] I. Iacopini, G. Petri, A. Barrat, and V. Latora, Simplicial models of social contagion, *Nature communications* **10**, 2485 (2019).
- [42] G. F. de Arruda, G. Petri, and Y. Moreno, Social contagion models on hypergraphs, *Physical Review Research* **2**, 023032 (2020).
- [43] F. Battiston, E. Amico, A. Barrat, G. Bianconi, G. Ferraz de Arruda, B. Franceschiello, I. Iacopini, S. Kéfi, V. Latora, Y. Moreno, *et al.*, The physics of higher-order interactions in complex systems, *Nature Physics* **17**, 1093 (2021).
- [44] G. Ferraz de Arruda, A. Aleta, and Y. Moreno, Contagion dynamics on higher-order networks, *Nature Reviews Physics* **6**, 468 (2024).
- [45] D. Lazer, K. Ognyanova, M. Baum, J. Druckman, J. Green, A. Gitomer, M. Simonson, R. H. Perlis, M. Santillana, A. Quintana, *et al.*, The COVID States Project# 43: COVID-19 vaccine rates and attitudes among Americans (2021).
- [46] D. Mistry, M. Litvinova, A. Pastore y Piontti, M. Chinazzi, L. Fumanelli, M. F. Gomes, S. A. Haque, Q.-H. Liu, K. Mu, X. Xiong, *et al.*, Inferring high-resolution human mixing patterns for disease modeling, *Nature communications* **12**, 323 (2021).
- [47] A. Aleta, G. Ferraz de Arruda, and Y. Moreno, Data-driven contact structures: from homogeneous mixing to multilayer networks, *PLoS computational biology* **16**, e1008035 (2020).
- [48] A. Cori, A.-J. Valleron, F. Carrat, G. S. Tomba, G. Thomas, and P.-Y. Boëlle, Estimating influenza latency and infectious period durations using viral excretion data, *Epidemics* **4**, 132 (2012).
- [49] M. Biggerstaff, S. Cauchemez, C. Reed, M. Gambhir, and L. Finelli, Estimates of the reproduction number for seasonal, pandemic, and zoonotic influenza: a systematic review of the literature, *BMC infectious diseases* **14**, 1 (2014).
- [50] K. E. Nelson and C. M. Williams, *Infectious disease epidemiology: theory and practice* (Jones & Bartlett Publishers, 2014).
- [51] O. W. in Data, Daily covid-19 vaccine doses administered, <https://ourworldindata.org/grapher/daily-covid-19-vaccination-doses?tab=chart&country=~USA> (2023), accessed on 2023-10-26.
- [52] U. Facts, What's the nation's progress on vaccinations?, <https://usafacts.org/visualizations/covid-vaccine-tracker-states/> (2023), accessed on 2023-10-26.
- [53] O. W. in Data, Daily share of the population receiving a covid-19 vaccine dose, <https://ourworldindata.org/grapher/daily-covid-19-vaccination-doses?tab=chart&country=~USA> (2023), accessed on 2023-11-2.
- [54] D. X. Morales, T. F. Beltran, and S. A. Morales, Gender, socioeconomic status, and covid-19 vaccine hesitancy in the us: an intersectionality approach, *Sociology of health & illness* **44**, 953 (2022).
- [55] E. Dubé and N. E. MacDonald, Vaccine acceptance: barriers, perceived risks, benefits, and irrational beliefs, in *The vaccine book* (Elsevier, 2016) pp. 507–528.
- [56] L. G. Alvarez-Zuzek, C. M. Zipfel, and S. Bansal, Spatial clustering in vaccination hesitancy: The role of social influence and social selection, *PLoS Computational Biology* **18**, e1010437 (2022).
- [57] M.-G. Hâncean, J. Lerner, M. Perc, J. L. Molina, and M. Geantă, Assortative mixing of opinions about covid-19 vaccination in personal networks, *Scientific reports* **14**, 3385 (2024).
- [58] X. Yuan and A. Crooks, From cyber space opinion leaders and the diffusion of anti-vaccine extremism to physical space disease outbreaks, in *Social, Cultural, and Behavioral Modeling: 10th International Conference, SBP-BRIMS 2017, Washington, DC, USA, July 5-8, 2017, Proceedings 10* (Springer, 2017) pp. 114–119.
- [59] S. L. Wilson and C. Wiysonge, Social media and vaccine hesitancy, *BMJ global health* **5**, e004206 (2020).
- [60] G. Jain, A. B. Sreenivas, S. Gupta, and A. A. Tiwari, The dynamics of online opinion formation: polarization around the vaccine development for covid-19, in *Causes and symptoms of socio-cultural polarization: role of information and communication technologies* (Springer, 2022) pp. 51–72.
- [61] S. Rathje, J. K. He, J. Roozenbeek, J. J. Van Bavel, and S. van der Linden, Social media behavior is associated with vaccine hesitancy, *PNAS nexus* **1**, pgac207 (2022).
- [62] A. L. Schmidt, F. Zollo, A. Scala, C. Betsch, and W. Quattrocioni, Polarization of the vaccination debate on facebook, *Vaccine* **36**, 3606 (2018).
- [63] W. Jennings, G. Stoker, H. Willis, V. Valgardsson, J. Gaskell, D. Devine, L. McKay, and M. C. Mills, Lack of trust and social media echo chambers predict covid-19 vaccine hesitancy, *MedRxiv*, 2021 (2021).
- [64] J. Müller, A. Tellier, and M. Kurschilgen, Echo chambers and opinion dynamics explain the occurrence of vaccination hesitancy, *Royal Society Open Science* **9**, 220367 (2022).
- [65] MOBS Lab, Mixing patterns repository, <https://github.com/mobs-lab/mixing-patterns> (2021), data and code repository associated with Mistry et al., *Nature Communications* (2021). Accessed: 2025-03-29.
- [66] U. C. Bureau, Population, U.S. Census Bureau, Population Division (2024), available at: <https://www.census.gov/topics/population.html>.
- [67] S. Arregui, A. Aleta, J. Sanz, and Y. Moreno, Projecting social contact matrices to different demographic structures, *PLoS computational biology* **14**, e1006638 (2018).
- [68] It can be shown that as $\langle k \rangle$ grows (while preserving the product $\beta \langle k \rangle$), the simulation outcomes approach the homogeneous mixing assumption, and then $r_{ER}(\infty) \rightarrow r_{hom}(\infty)$.
- [69] Y. Moreno, R. Pastor-Satorras, and A. Vespignani, Epidemic outbreaks in complex heterogeneous networks, *The European Physical Journal B-Condensed Matter and Complex Systems* **26**, 521 (2002).

[70] It should be noted that $v(0)$ represents an initial condition in a population otherwise naive to the disease. In reality, however, the “already vaccinated” category in

surveys preceded the second wave of COVID-19 in the US.