



# Modeling the Interplay Between Seasonal Flu Outcomes and Individual Vaccination Decisions

Irena Papst<sup>1</sup>  · Kevin P. O’Keeffe<sup>2</sup>  · Steven H. Strogatz<sup>3</sup> 

Received: 19 January 2021 / Accepted: 28 December 2021 / Published online: 31 January 2022  
© The Author(s), under exclusive licence to Society for Mathematical Biology 2022

## Abstract

Seasonal influenza presents an ongoing challenge to public health. The rapid evolution of the flu virus necessitates annual vaccination campaigns, but the decision to get vaccinated or not in a given year is largely voluntary, at least in the USA, and many people decide against it. In some early attempts to model these yearly flu vaccine decisions, it was often assumed that individuals behave rationally, and do so with perfect information—assumptions that allowed the techniques of classical economics and game theory to be applied. However, these assumptions are not fully supported by the emerging empirical evidence about human decision-making behavior in this context. We develop a simple model of coupled disease spread and vaccination dynamics that instead incorporates experimental observations from social psychology to model annual vaccine decision-making more realistically. We investigate population-level effects of these new decision-making assumptions, with the goal of understanding whether the population can self-organize into a state of herd immunity, and if so, under what conditions. Our model agrees with the established results while also revealing more subtle population-level behavior, including biennial oscillations about the herd immunity threshold.

**Keywords** Seasonal influenza · Vaccination · Decision-making · Social psychology · SIR model

## 1 Introduction

Annual influenza epidemics are a significant public health challenge, with up to 650,000 individuals dying from respiratory diseases associated with the flu each year

---

✉ Irena Papst  
ip98@cornell.edu

<sup>1</sup> Center for Applied Mathematics, Cornell University, Ithaca, NY, USA

<sup>2</sup> Senseable City Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

<sup>3</sup> Department of Mathematics, Cornell University, Ithaca, NY, USA

(World Health Organization 2017). In the USA alone, the total economic burden of seasonal influenza, including direct medical costs and lost earnings due to illness or death, has been estimated as \$26.8 billion annually (Molinari et al. 2007).

One of the main challenges of controlling seasonal influenza spread is that the viruses evolve quickly (on the same time scale as the annual epidemics), with multiple strains circulating concurrently. A key adaptation mechanism, antigenic drift, gives rise to new influenza strains by randomly changing segments of viral surface proteins. Given that a host's immune system uses these surface proteins to identify the virus so that it may be neutralized (Taubenberger and Kash 2010), antigenic drift thus acts as an evolutionary countermeasure. It helps the flu evade the immune system and thereby promotes its spread through the host population.

The rapid evolution of the flu results in the constant threat of a pandemic, and it also makes it challenging to develop effective, long-lasting vaccines. The seasonal flu vaccine is updated every year to protect against the strains that seem to pose the largest upcoming threat. Seasonal influenza vaccination is largely voluntary in the USA, so individuals must decide whether or not to vaccinate each year.

For many people, this decision is not easy. It involves many quantities that are effectively impossible for an individual to estimate accurately, such as their likelihood of being vaccinated successfully, the probability of an adverse reaction to the vaccine, as well as their increased risk of catching the flu by foregoing the vaccine. Yet although the decision may be difficult at an individual level, at a societal level the benefits of vaccinating can be immense; if a critical mass of individuals choose to immunize themselves, "herd immunity" can be achieved. In this desirable state, the density of susceptible individuals is so low that an infection chain cannot be sustained and so an epidemic cannot occur (Fine 1993).

Somewhat paradoxically, the possibility of achieving herd immunity makes an individual's decision to vaccinate or not even more complex, at least when viewed through the lens of classical game theory (wherein agents are assumed to be purely self-interested, and to behave rationally with perfect information about the situation at hand). The issue is that as vaccination coverage increases, individuals are increasingly incentivized *not* to vaccinate. Each person would do better relying on others to bear any burden associated with vaccination, while everyone reaps the benefits of widespread immunity. Such "free-riding" logic makes it impossible to ever actually achieve herd immunity with rational hosts.

The free-riding problem is common in (though not exclusive to) models where agents rationally weigh a delayed collective group benefit against immediate individual costs.<sup>1</sup> Individuals are assumed to make decisions by selecting the strategy that maximizes an objective, individual payoff function, as prescribed by classical economics (Hardin 2013). Some early models of voluntary vaccination decisions involve such assumptions and inevitably yield agents that utilize free-riding logic, which precludes herd immunity (Geoffard and Philipson 1997; Bauch and Earn 2004).

However, a more recent empirical study suggests that free-riding logic is uncommon when individuals specifically consider whether or not to get the seasonal influenza

---

<sup>1</sup> Free-riding occurs when individuals are less likely to get vaccinated if they observe reduced disease prevalence, whether or not their vaccination decision is made "rationally."

vaccine (Parker et al. 2013). In fact, this study finds that the majority of individuals surveyed do not account for the vaccination decisions of others when making their own decision. Nevertheless, despite increasing evidence that assumptions from classical economics may not appropriately capture human behavior in the context of infectious disease spread, some behavior-disease models continue to be built on such foundations (Verelst et al. 2016).

Other studies have challenged these assumptions by replacing them with those from behavioral economics, which leverages social psychology in its models of human decision-making. Voinson et al. (2015) develop a behavior-disease model that incorporates cognitive biases and differing vaccine opinions among individuals to study vaccination coverage over time. Oraby and Bauch (2015) study pediatric acceptance of vaccines by incorporating prospect theory into their disease model.

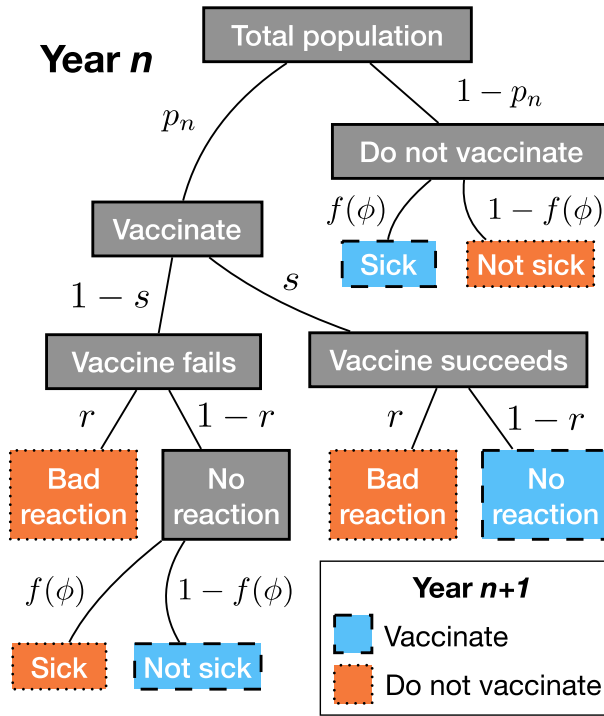
In this paper, we consider a simplified model for the interplay between annual vaccination decisions and seasonal influenza spread, in which individual voluntary vaccination decisions are informed by observed social psychology in this context. Unlike previous work, we model *repeated* vaccination decisions to reflect the annual vaccination decision necessitated by the rapid evolution of influenza viruses. We investigate population-level effects of these new decision-making assumptions, with the goal of understanding whether the population can self-organize into a state of herd immunity, and if so, under what conditions. Despite our model's idealized nature, we find that its results align with those utilizing assumptions based on classical economics, although our model also predicts more nuanced population-level behaviors, such as oscillations in and out of herd immunity on a biennial basis.

## 2 Conceptual Model

Decision theory and social psychology suggest that, in general, individuals tend to use heuristics, or rules of thumb, rather than a "rational" cost-benefit analysis in complex decision-making (Tversky and Kahneman 1974). Moreover, decision-making tends to obey the law of inertia: choices generally remain unchanged but are sensitive to both small nudges and unfavorable resulting outcomes (Thaler and Sunstein 2009). Our model is based on both of these ideas.

For simplicity, assume that each year, an individual chooses whether or not to receive the seasonal influenza vaccine based solely on evaluating their most recent outcome with both the vaccine and the disease. The vaccine carries a cost, or risk, of adverse reaction, which can be interpreted as a cost to an individual's health due to vaccine side effects (morbidity), a direct economic cost from paying for the vaccine, and/or an indirect economic cost such as taking unpaid leave from work to get vaccinated. In what follows, we will interpret vaccine cost as morbidity, but the modeling framework is flexible enough to accommodate other interpretations. There is also some probability that vaccination successfully confers immunity upon the recipient.

Before we present the model in detail in Sect. 3, let us first describe it intuitively. In any given year, individuals make the decision of whether or not to vaccinate, then follow through with their choice, and then the flu season occurs. The epidemic resolves itself, each individual assesses their personal outcome from the past year and then



**Fig. 1** All possible decisions and outcomes for an individual in year  $n$ , leading to different decisions for year  $n + 1$ . Boxes with solid borders denote decisions and possible intermediate repercussions in year  $n$ . Boxes with dashed and dotted borders denote final repercussions that determine the decision in year  $n + 1$ : vaccinate (box with dashed border) or do not vaccinate (box with dotted border). In the model,  $p_n$  is the proportion of vaccinators in year  $n$ ,  $r$  is the probability that the vaccine induces a cost,  $s$  is the probability that the vaccine succeeds,  $\phi$  is the final size of the epidemic (normalized as a fraction of the whole population), and  $f(\phi) = \phi(sp_n)/(1 - sp_n)$  is the fraction of the susceptible population that was infected in the epidemic occurring in year  $n$ . A “bad reaction” results when one incurs a cost from the vaccine (e.g., vaccine side effects on health, or economic burden). See Sect. 3 for model details

decides whether or not to get vaccinated prior to next year’s flu season. The decision rule is a simple heuristic: If a person “won” last year (did not get sick and did not have an adverse reaction to the vaccine), they stick to their vaccination choice and make the same decision the following year. If they “lost” (got sick or had a bad reaction to the vaccine), they are nudged to change their behavior (switch from vaccinating to not, or vice versa) in hopes of eliciting a better outcome during the next flu season.

Figure 1 shows a schematic of the model. As an example of how to flow through the chart, let us first consider an individual who has decided to vaccinate in a given year (a decision made by a proportion  $p_n$  of the population in year  $n$ ). This decision corresponds to the top left fork of the tree in Fig. 1. Next, suppose the vaccine succeeds in conferring immunity (which occurs with probability  $s$ ), and the vaccine does not elicit an adverse reaction (which occurs with probability  $1 - r$ ). This favorable outcome does not push the individual away from their default (winning) strategy of vaccinating, so they decide to vaccinate again the following year; the vaccine seems

to have succeeded in protecting them from the flu. (We are assuming here that only the failure of the vaccine can be observed by the individual, and only if they happen to get ill that year. Otherwise, the success of the vaccine is presumed, since there is no evidence to the contrary.)

On the other hand, an individual may choose *not* to vaccinate in a given year (a decision made by a proportion  $1 - p_n$  of the population, shown by following the top right fork in Fig. 1). If such an individual then happens to contract the flu (which occurs with a probability that we will calculate below), this non-vaccinating individual's choice was a "losing" strategy, and so they decide to vaccinate the following year. Various other paths through the tree can be understood similarly.

In the next section, we write down the governing equations for our model. We begin by recalling some standard results for a simple epidemiological model and then couple that model to a social psychological model for individual vaccination decisions.

### 3 Mathematical Model

#### 3.1 Epidemic Model

We choose the susceptible–infected–removed (SIR) model of infectious disease dynamics (Kermack and McKendrick 1927), both because it is a well-established epidemiological model and because there exists an analytical expression for the final size of epidemics predicted by the model. The SIR model is adequate for modeling each annual influenza outbreak individually, though we note that a more realistic flu model could be substituted into our vaccination coverage model, provided that the final size could at least be approximated numerically.

Suppose  $S(t)$  and  $I(t)$  represent proportions of a population that are susceptible and infected, respectively, at a time  $t$ . Infected individuals are assumed to be immediately infectious; there is no latency period in this model. There may also be individuals that are removed from the infection process as they have already recovered from the illness (denoted by proportion  $R(t)$ ), but since we assume this disease propagates in a closed population, we have  $R(t) = 1 - S(t) - I(t)$ , which means we do not need to track the removed individuals explicitly.

The SIR model is defined by a set of two coupled, nonlinear, ordinary differential equations:

$$\frac{dS}{dt} = -\beta SI, \quad (1a)$$

$$\frac{dI}{dt} = \beta SI - \gamma I, \quad (1b)$$

where  $\beta$  is the disease transmission rate and  $\gamma$  is the disease recovery rate. The derived quantity  $\mathcal{R}_0 = \beta/\gamma$  is the basic reproduction number of the disease; it gives the average number of secondary cases generated by an infectious individual in a fully susceptible population over the course of their illness.

In our model, we assume the initial conditions

$$S(0) = 1 - sp, \quad 0 < I(0) \ll 1, \quad R(0) = sp \quad (2)$$

to incorporate a vaccine uptake level at proportion  $p$  with vaccine efficacy  $s$ .

### 3.2 Final Size of the Epidemic

For this model, one can derive an implicit equation for the final size of the epidemic  $\phi(x)$ , where  $x$  is the proportion initially immune to the disease (Ma and Earn 2006):

$$\phi(x) = (1 - x)(1 - e^{-\mathcal{R}_0\phi(x)}). \quad (3)$$

The solution to Eq. 3 can be written in terms of the principal branch of the product log function (i.e., the Lambert W-function), denoted by  $W[\cdot]$ :

$$\phi(x) = 1 - x + \frac{1}{\mathcal{R}_0} W \left[ -\mathcal{R}_0(1 - x)e^{-\mathcal{R}_0(1-x)} \right] \quad (4)$$

This expression for the final size of the yearly influenza epidemic is used later in Eq. 7 to complete the model.

### 3.3 Critical Vaccination Threshold for Herd Immunity

An epidemic cannot be sustained if the average number of secondary cases provoked by an infected individual in the population is below one. (Since this infected individual cannot even replace themselves in the infection chain, let alone generate further infections.) In other words, for the population to achieve herd immunity, the *effective* reproduction number (the basic reproduction number times the proportion currently susceptible),  $\mathcal{R}_{\text{eff}} = \mathcal{R}_0(1 - sp)$ , must be driven below 1. Thus the critical vaccination threshold,  $p_{\text{crit}}$ , satisfies the equation  $\mathcal{R}_0(1 - sp_{\text{crit}}) = 1$ , and so

$$p_{\text{crit}} = \frac{1}{s} \left( 1 - \frac{1}{\mathcal{R}_0} \right). \quad (5)$$

### 3.4 Estimating $\mathcal{R}_0$ for Seasonal Influenza

Estimates of  $\mathcal{R}_0$  vary depending on year, location, and influenza subtype since the basic reproduction number depends not only on the immunological properties of the virus, but also on the social behavior of the host population. A systematic review by Biggerstaff et al. (2014) catalogues many estimates of both the basic and effective reproduction numbers of pandemic, zoonotic, and seasonal influenza.

The most relevant estimates for our study are those for the 1976–1981 outbreak of H1N1/H3N2/B in the USA. Two studies were performed to estimate the basic reproduction number in this outbreak, and they both use serologically confirmed infections for their data, which make these estimates particularly reliable.<sup>2</sup> One study found  $\mathcal{R}_0 = 1.70$  (Ferguson et al. 2006), while another found  $\mathcal{R}_0 = 1.16$  (Britton and Becker 2000) for this outbreak. We average these two values and take  $\mathcal{R}_0 = 1.4$  as a reasonable estimate of the basic reproduction number for seasonal influenza in a modern US population. While the model results presented in the main text all use  $\mathcal{R}_0 = 1.4$ , we have explored a range of  $\mathcal{R}_0$  values in Appendix B (Figures S1–S2).

### 3.5 Vaccination Decision Model

Let  $p_n$  be the proportion of the population that vaccinates in year  $n$ . Our goal in this section is to derive a discrete map for the vaccine coverage  $p_{n+1}$  in year  $n + 1$ .

In the following argument, we assume the vaccine is fully quantified by its cost and its efficacy. For the purposes of this discussion, we will think of the cost in terms of vaccine morbidity (side effects to immunization), though the cost could be interpreted as an economic one (for instance, if individuals have to pay for the vaccine or take unpaid time off of work to obtain it). We denote the cost, or probability of vaccine morbidity, by  $0 \leq r \leq 1$ , and the vaccine efficacy by  $0 \leq s \leq 1$ .

To ease the notation in the derivation below, it proves useful to introduce a function  $f(\phi)$  to denote the proportion of all *susceptible* individuals who get sick during an epidemic of size  $\phi$ . To calculate  $f$  in terms of  $\phi$ , note that the fraction of the total population that is susceptible in year  $n$  is  $1 - sp_n$ . Of these individuals, a fraction  $f \cdot (1 - sp_n)$  will get infected, by definition of  $f$ . But since this fraction also equals the number of infected individuals divided by the total population, it simply equals  $\phi$ , the fractional size of the epidemic, as given by Eq. 4. Therefore,  $\phi = f \cdot (1 - sp_n)$ , from which we conclude that

$$f(\phi(sp_n)) = \frac{\phi(sp_n)}{1 - sp_n}. \quad (6)$$

In other words, the proportion of susceptible individuals who end up infected is simply the final size of the epidemic renormalized to the susceptible population.

With these preliminaries out of the way, we can deduce the vaccine coverage rate  $p_{n+1}$  in year  $n + 1$  by considering all possible outcomes for an individual based on their choice of whether or not to vaccinate in year  $n$  and by counting the proportion of the population flowing down each of the branches in Fig. 1 into vaccinating in year  $n + 1$ . We assume that every individual is susceptible to that year's flu strain at the start of each flu season, so everyone must make the choice of whether or not to vaccinate each year.

First consider the group of non-vaccinating individuals, which make up a proportion  $1 - p_n$  of the population. These individuals will only vaccinate in year  $n + 1$  if they

<sup>2</sup> It is difficult to distinguish seasonal influenza from other upper respiratory tract infections by symptoms alone, so studies based purely on reported symptoms may not yield a good estimate for the  $\mathcal{R}_0$  of seasonal influenza. Instead, studies based on serologically confirmed infections are more reliable.

get sick in year  $n$ , an event that occurs to a fraction  $f = f(\phi(sp_n))$  of them. Thus, the equation for  $p_{n+1}$  will include a term  $f \cdot (1 - p_n)$ , which accounts for those that did not vaccinate and got sick.

For vaccinating individuals, either the vaccine succeeds, with probability  $s$ , or it does not, with probability  $1 - s$ . If the vaccine succeeds, there is still an independent chance that the individual will have side effects that discourage them from vaccinating the following year, which occurs at the vaccine morbidity rate,  $r$ . However, those for whom the vaccine successfully conferred immunity and provoked no side effects will once again vaccinate the following year since they have no reason to change strategy, which adds the term  $(1 - r) \cdot s \cdot p_n$  to the equation for  $p_{n+1}$ .

If the vaccine fails for an individual (with probability  $1 - s$ ), but did not cause any discouraging side effects (with probability  $1 - r$ ), the only reason they would continue to vaccinate would be if they thought the vaccine succeeded; that is, they happened not to get sick, even though they were not successfully immunized. A proportion  $1 - f$  of susceptible individuals avoid infection, so the final term of the equation for  $p_{n+1}$  is  $(1 - f) \cdot (1 - r) \cdot (1 - s) \cdot p_n$ .

Putting all of these contributions together, we find that the discrete map for  $p_{n+1}$ , the proportion of the population vaccinating in year  $n + 1$ , is given by

$$p_{n+1} = \frac{\phi(sp_n)}{1 - sp_n}(1 - p_n) + (1 - r)sp_n + \left[1 - \frac{\phi(sp_n)}{1 - sp_n}\right](1 - r)(1 - s)p_n, \quad (7)$$

where the function  $\phi$  is given by Eq. 4. This map is biologically sensible; if  $0 \leq p_n \leq 1$ , one can check that  $0 \leq p_{n+1} \leq 1$ . Hence, as long as the initial condition is sensible ( $0 \leq p_0 \leq 1$ ), all subsequent iterations remain in  $[0, 1]$ .

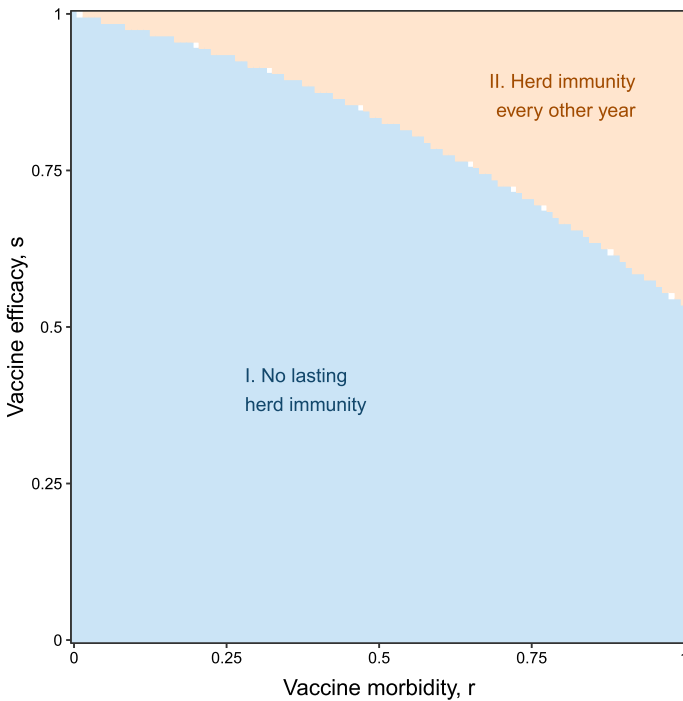
## 4 Results

### 4.1 Model Predictions

The predictions of the model depend on the relative magnitudes of its parameters: the vaccine parameters (morbidity or cost,  $r$ , and efficacy,  $s$ ) and the disease parameter,  $\mathcal{R}_0$ . The basic reproduction number  $\mathcal{R}_0$  gives a sense of the “infectiousness” of the disease; in our analysis, we estimate the basic reproduction number of seasonal influenza in a modern US population to be  $\mathcal{R}_0 = 1.4$  (see Sect. 3 for details), indicating that a person infected with seasonal influenza will infect on average 1.4 other people in a fully susceptible population. For the vaccine parameters, we note that seasonal flu vaccines have very low morbidity (Centers for Disease Control 2016), and their efficacy varies from low to moderate (Osterholm et al. 2012; World Health Organization 2015; Centers for Disease Control 2017).

The best case would be for the population to self-organize into herd immunity, by driving the proportion of the population vaccinated above the critical threshold,  $p_{\text{crit}}$ . When vaccine coverage meets or exceeds this threshold, no epidemic occurs. One might expect the model to self-organize into herd immunity if the population can collectively make use of the memory of the previous flu seasons in a lasting way.







**Fig. 2** Long-term model behavior for  $\mathcal{R}_0 = 1.4$  and  $p_0 = 0$ , as a function of vaccine morbidity ( $0 < r \leq 1$ ) and vaccine efficacy ( $0 \leq s \leq 1$ ). Behavior in these regions of parameter space was deduced by iterating the vaccine coverage map (Eq. 7) numerically until it converged to a fixed point, i.e., until the discrepancy between successive iterations fell below a tolerance of  $10^{-8}$ . (White points are parameter values for which the simulation was not able to converge to a solution before the maximum number of iterations was reached.) The majority of parameter space is dominated by convergence to vaccine levels below the herd immunity threshold, which results in no lasting herd immunity (region I: the system converges to a period 1 fixed point,  $p^*$ , that satisfies  $p^* < p_{crit}$ ). For higher vaccine efficacy, there is a possibility of achieving herd immunity every other year, provided that vaccine morbidity  $r$  is sufficiently large (region II: the system converges to a period-2 fixed point,  $(p_1^*, p_2^*)$ ). In this regime, the system oscillates between sub-optimal vaccine coverage ( $p_1^* < p_{crit}$ ) and herd immunity with overvaccination ( $p_2^* > p_{crit}$ ): see Fig. 3A

To our disappointment, we find that if there is any cost to the vaccine ( $r > 0$ ), our model cannot self-organize into lasting herd immunity. There are two main regions of parameter space in this case (Fig. 2): a large region where the system eventually converges to vaccine levels below the herd immunity threshold (region I), and a smaller region where the system oscillates in and out of the herd immunity region on a yearly basis (region II).

Note that where game-theoretic models always predict free-riding that make herd immunity impossible to achieve in the long term, our model predicts that it is possible for the population to achieve herd immunity every other year, even if there is a cost and moderate failure rate to the vaccine. This oscillatory behavior (Fig. 3A) is the result of the system converging to a state where it alternates between the population bearing a significant disease burden (a large epidemic in the previous year encouraging vaccination in the following year) and the population bearing a significant vaccine cost



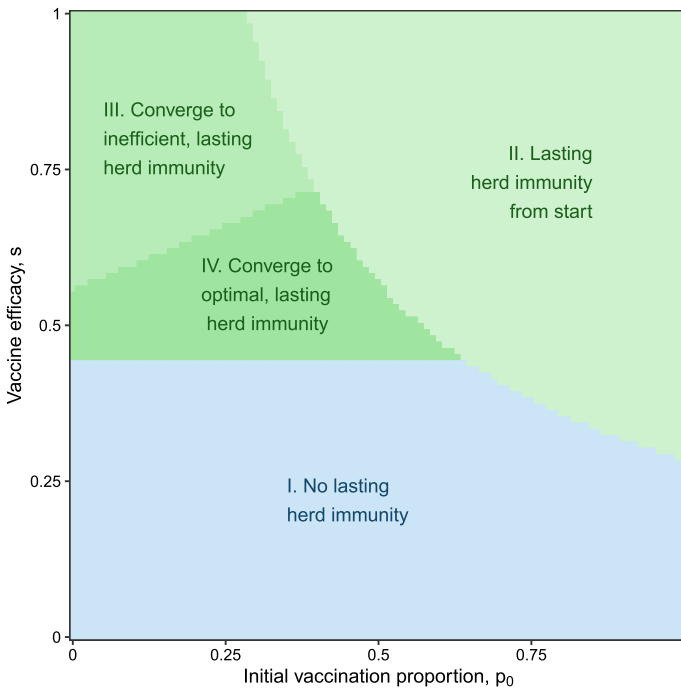
**Fig. 3** Vaccine coverage level in the regime where herd immunity eventually occurs every other year ( $\mathcal{R}_0 = 1.4, r = 0.8, s = 0.7$ ). Panel **A** Vaccine coverage level over time. The system converges to a state where the vaccine coverage level oscillates asymmetrically about the critical vaccination threshold,  $p = p_{crit}$ , denoted by the dashed line. Panel **B** Dependence of the next year's vaccine coverage level,  $p_{n+1}$ , on that of the current year,  $p_n$ , for this parameter set



**Fig. 4** Number of years spent in the herd immunity interval ( $p > p_{crit}$ ) during the transient period in the regime of no lasting herd immunity ( $\mathcal{R}_0 = 1.4, p_0 = 0$ ). As the vaccine improves in quality (either vaccine morbidity decreases, or vaccine efficacy increases), the time period spent in the herd immunity interval lengthens. The dependence of the number of years spent in the herd immunity interval on the initial condition,  $p_0$ , is explored in Figure S3 (Appendix B)

(which incentivizes non-vaccination *en masse* in the following year). When a vaccine has a moderate-to-high efficacy, and a sufficiently high cost, the system is constantly balancing an illness-vaccine cost trade off due to the built-in short term memory of individuals in the model.

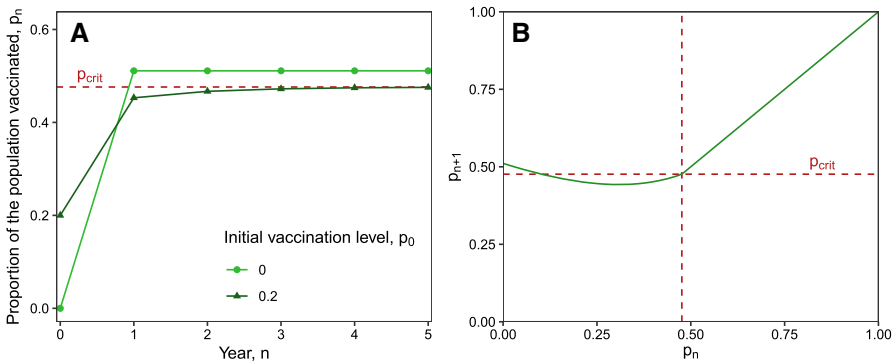
Even in the case where there is no *lasting* herd immunity, the system may nevertheless spend a significant length of time in the herd immunity interval (Fig. 4) before eventually dropping out and converging to a stable point that is below the herd immunity threshold. This effect is especially pronounced when vaccine morbidity is low, and even when the vaccine is only moderately efficacious, both of which are properties of the real seasonal influenza vaccines. The transient herd immunity period increases as vaccine efficacy increases and/or vaccine cost decreases.



**Fig. 5** Long-term model behavior with no vaccine cost ( $r = 0$ ) for  $\mathcal{R}_0 = 1.4$ . Behavior in these regions of parameter space was deduced using the analytical criteria detailed in Appendix A. There is still a region with no lasting herd immunity (region I: the system converges to a fixed point,  $p^*$ , that satisfies  $p^* < p_{crit}$ ); in this regime, the system never achieves herd immunity ( $p_n < p_{crit}$  for all  $n \geq 0$ ). However, the system exhibits self-organized herd immunity when vaccine efficacy is sufficiently high, through a variety of mechanisms. The system may start (and therefore stay) in the herd immunity interval (region II:  $p_0 \geq p_{crit}$ ), it may converge to “inefficient” lasting herd immunity (region III: sustained overvaccination), or it may converge to “optimal” lasting herd immunity (region IV: vaccination approaching the herd immunity threshold  $p_{crit}$ )

When there is no cost to the vaccine ( $r = 0$ ; Fig. 5), the system can self-organize into herd immunity in three ways (regions whose label contains “lasting herd immunity”), in addition to yielding no lasting herd immunity as before (region I). The system may start in the herd immunity region and therefore stay in it indefinitely (region II), since there is no vaccine cost to drive the coverage level down. Alternatively, the system may converge to lasting herd immunity which is either inefficient as it involves overvaccination (region III), or it may converge to optimal, lasting herd immunity precisely at the herd immunity threshold (region IV).

The mechanisms that drive the population to either inefficient or optimal self-organized herd immunity are markedly different (Fig. 6A). In the case of sustained overvaccination, the population starts at a relatively low level of vaccination initially (lighter curve). A substantial epidemic occurs in the first year, encouraging a large proportion of individuals to vaccinate in the following year: too many, in fact. The population springs itself into the herd immunity interval after that first year, and since there is no opposing force pushing vaccination coverage down, overvaccination continues indefinitely. In the case of optimal vaccination, the population may also start



**Fig. 6** Vaccine coverage level with no vaccine cost ( $r = 0$ ) in the regime of self-organized herd immunity ( $\mathcal{R}_0 = 1.4$ ,  $s = 0.6$ ). Panel **A** Vaccine coverage level over time. If the initial population level is too low (lighter curve), the population springs into the interior of the herd immunity interval,  $[p_{crit}, 1]$ , which results in sustained overvaccination. If the population initially vaccinates at a more moderate level (darker curve), vaccine coverage converges to the optimal herd immunity threshold,  $p_{crit}$ , in an asymptotic way. Panel **B** Dependence of the next year's vaccine coverage level,  $p_{n+1}$ , on that of the current year,  $p_n$ , for this parameter set

at a relatively low level of initial vaccination, but the first epidemic sustained is not as devastating as in the previous case (darker curve). A moderate proportion of the population is affected by the disease and switches to vaccinating in the following year. The epidemic sustained in this next year is not as large as the one before it (thanks to the increase in vaccination) and encourages another (smaller) group of individuals to switch to vaccinating next year. This process gradually guides the population to the herd immunity threshold, eventually achieving the optimal level of vaccination.

## 5 Discussion

While the simple model studied here cannot self-organize into sustained herd immunity when there is any cost to the vaccine, it may still achieve herd immunity every other year. When the vaccine efficacy is sufficiently large for a given cost, there is an ongoing battle between the disease and the vaccine. If the population undervaccinates in one year, it undergoes an epidemic which drives the system to overvaccinate in the following year, thanks to the short-term memory of individuals in the model. A non-trivial proportion of the population then bears some cost associated with the vaccine, which discourages those individuals from getting vaccinated the following year, driving the system back down to an undervaccinated state, and the cycle repeats. Similar biennial behavior driven by individual short-term memory was previously reported in another behavior-disease model (Wells and Bauch 2012), though only in the transient dynamics. In our case, the biennial cycle occurs asymptotically.

Provided the amplitude of this (asymmetric) oscillation about the herd immunity threshold is sufficiently small, this regime effectively achieves herd immunity as any epidemic that occurs is relatively small. While the goal of disease eradication has not strictly been achieved, the resulting epidemics are so small in the model that a bit

of stochasticity may be enough to push the circulating flu strain into extinction (in a closed population).

Although this promising biannual behavior is possible, the region of no lasting herd immunity dominates vaccine parameter space, particularly for vaccine morbidity and efficacies that are realistic for the seasonal influenza vaccine (morbidity near zero, efficacy around 50%) (Fig. 2). The system may not drive itself to herd immunity asymptotically for these types of vaccines, but a significant length of time is spent in the herd immunity interval during the transient period (Fig. 4). This effect opens the door for other public health interventions (e.g., vaccination- and disease-awareness campaigns) which have not been included in the model but may help push the population into lasting herd immunity. Increases to the length of time spent in the herd immunity interval can be achieved by improving the vaccine, by increasing vaccine efficacy, and/or by decreasing vaccine morbidity.

In the case where there is no morbidity to the vaccine, the system can achieve self-organized herd immunity that is either inefficient (due to overvaccination) or optimal (at the critical vaccination threshold). If the population initially vaccinates at a very low level, it undergoes a large epidemic, and the following year, the population overreacts, propelling itself into the herd immunity interval much like a diver on a springboard. Overvaccination continues since there is no cost to the vaccine, and thus, no force pushing population vaccine coverage down. On the other hand, moderate initial vaccination leads the population to converge to the optimal vaccination level at the herd immunity threshold. In this case, the population gradually learns from year to year through successively smaller epidemics. Each such epidemic recruits smaller and smaller proportions of the population to vaccinate until herd immunity is achieved. Notably, this result occurs even with a moderately effective vaccine, like that of real seasonal influenza. This result is particularly striking as it follows from individuals seeking to optimize their own outcomes, yet they also collectively achieve the social optimum.

While it may not be realistic to assume that a vaccine can be considered costless to an individual, this extreme case illustrates that if a vaccine can be perceived as costless, the population can self-organize into sustained herd immunity. Such a result is still possible even if the vaccine is only moderately efficacious and even if the population does not immediately take to getting vaccinated.

## 6 Conclusion

We have presented an intentionally simple model for seasonal influenza vaccination that makes use of the established results in social psychology to inform individual vaccine decision-making, giving rise to both interesting and interpretable dynamics.

In the case where there is some cost to the vaccine, our model still predicts regimes where vaccination coverage is below the herd immunity threshold in the long term, which agrees with previous models. However, our model also predicts new regions where herd immunity is achieved every other year: a result of the population oscillating between vaccine-based and disease-based morbidity. When we further assume that the vaccine has no cost, it is still possible for the model to predict no lasting herd immunity. We also observe convergence to enduring herd immunity, either at the

optimal level, where the population vaccinates exactly enough to reach this protected state, or inefficiently, where the population overvaccinates.

Our disease-behavior model is deliberately simple as a first step, to focus on the effect of incorporating a more realistic decision model on top of a well-established model for disease spread. Future work should focus on making this model more realistic and validating it with appropriate data.

Currently, both the decision-making and disease processes are deterministic; a stochastic version of this model in either respect would be closer to reality. Since agents only rely on the current state of the system to inform their next decision, our model could easily be cast in a Markov chain framework. Modeling the disease spread on a socio-spatial network would also provide greater realism, by mimicking the way hosts interact and thus spread infectious diseases like the flu (Chao et al. 2010).

The vaccine decision model could also be extended to include social norms (Oraby et al. 2014), especially prosocial motives (Li et al. 2016), which may play a significant role in real-life vaccination decisions. Exposure to expert opinions on vaccination (Chapman and Coups 1999), such as those of primary care physicians, as well as non-expert opinions, like misinformation (Loomba et al. 2021), are also effects worth incorporating in future models, depending on the disease and social context.

Seasonal influenza is an immensely complex phenomenon, and we have not accounted for issues such as that of multiple concurrently circulating strains (Prosper et al. 2011), cross-reactivity of vaccines between strains (Iorio et al. 2012; Moa et al. 2016), naturally-acquired immunity carried over from previous flu seasons (Krammer 2019; Lee et al. 2019), or waning vaccine immunity over the course of the flu season (Rambhia and Rambhia 2018). Such additions to the model would also serve to make it more realistic.

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1007/s11538-021-00988-z>.

**Funding** IP was supported by a Natural Sciences and Engineering Research Council of Canada postgraduate scholarship (PGS-D). KPO and SHS were supported by NSF Grant Nos DMS-1513179 and CCF-1522054.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Code availability** All code required to generate the figures presented in this paper is available on Github (<https://github.com/papsti/vaxgame>).

## Appendix

### A. Analytical Criteria for Long-Term Behavior in Model with No Cost ( $r = 0$ )

Figure 5 can be produced by directly iterating the map to numerical convergence, but it can be produced equivalently using analytical criteria based on the existence and location of fixed points.

When  $r = 0$ , the model map (Eq. 7) reduces to

$$p_{n+1} = \frac{\phi(sp_n)}{1 - sp_n}(1 - p_n) + sp_n + \left[1 - \frac{\phi(sp_n)}{1 - sp_n}\right](1 - s)p_n. \tag{A 1}$$

Fixed points,  $p$ , of this map must satisfy

$$p = \frac{\phi(sp)}{1 - sp}(1 - p) + sp + \left[1 - \frac{\phi(sp)}{1 - sp}\right](1 - s)p, \tag{A 2}$$

which simplifies to

$$\frac{\phi(sp)}{1 - sp}(p(2 - s) - 1) = 0. \tag{A 3}$$

Equation A 3 is satisfied when either (i)  $\phi(sp) = 0$  or (ii)  $p(2 - s) - 1 = 0$ .

Case (i) is satisfied by any  $p \in [p_{\text{crit}}, 1]$  since  $\phi(sp) = 0$  if and only if  $p \geq p_{\text{crit}}$ . Provided  $p_{\text{crit}} < 1$ , the ‘‘herd immunity’’ interval  $[p_{\text{crit}}, 1]$  is an invariant set of neutrally stable fixed points that exists in the map’s domain of  $[0, 1]$ . In other words, if  $p_n \in [p_{\text{crit}}, 1]$  for any  $n$ , the trajectory is then trapped in the herd immunity interval for all remaining time (precisely at the value  $p_n$ ).

Region II in Fig. 5 is given by all  $(p_0, s)$  that satisfy  $p_0 \geq p_{\text{crit}} = \frac{1}{s} \left(1 - \frac{1}{\mathcal{R}_0}\right)$ . In other words, the population starts at herd immunity and remains at herd immunity indefinitely. Region III is given by all  $(p_0, s)$  such that  $p_0 < p_{\text{crit}}$  but  $p_1 \geq p_{\text{crit}}$ . This region represents populations whose first epidemic was so large that it propels the population into herd immunity immediately after the first year.

Case (ii) is satisfied by  $p^* = 1/(2 - s)$ . This fixed point is disjoint from the herd immunity interval when

$$p^* < p_{\text{crit}}, \tag{A 4}$$

$$\frac{1}{2 - s} < \frac{1}{s} \left(1 - \frac{1}{\mathcal{R}_0}\right), \tag{A 5}$$

$$s < 1 - \frac{1}{2\mathcal{R}_0 - 1}. \tag{A 6}$$

Let us define  $s_{\text{crit}} = 1 - 1/(2\mathcal{R}_0 - 1)$ ; when  $s \geq s_{\text{crit}}$ , the fixed point  $p^* = 1/(2 - s)$  disappears into the herd immunity interval.

Numerical simulations suggest that  $p^*$  is stable when it exists disjoint from the herd immunity interval (i.e., when  $s < s_{\text{crit}}$ ), with the basin of attraction being  $[0, p_{\text{crit}}]$ ; such trajectories will converge to  $p^*$  as  $n \rightarrow \infty$ . Region I is given by all  $(p_0, s)$  that satisfy  $p_0 < p_{\text{crit}}$  and  $s < s_{\text{crit}}$ . In other words, vaccine coverage starts below the herd immunity threshold and never surpasses it. Instead, vaccine coverage converges to  $p^* = 1/(2 - s) < p_{\text{crit}}$ .

Lastly, region IV is given by all  $(p_0, s)$  with both  $p_0 < p_{\text{crit}}$  and  $p_1 < p_{\text{crit}}$ , but also  $s \geq s_{\text{crit}}$ . In this case, the fixed point  $p = 1/(2 - s)$  does not exist distinct from

the herd immunity interval, but also the first epidemic is not strong enough propel vaccine coverage over the herd immunity threshold ( $p_1 < p_{\text{crit}}$ ). Under these conditions, numerical simulations suggest that  $\{p_0, p_1, \dots\}$  is a monotonically increasing sequence where each  $p_n < p_{\text{crit}}$ , but  $\lim_{n \rightarrow \infty} p_n = p_{\text{crit}}$ .

## References

- Bauch CT, Earn DJD (2004) Vaccination and the theory of games. *Proc Natl Acad Sci U S A* 101(36):391–394. <https://doi.org/10.1073/pnas.0403823101>
- Biggerstaff M, Cauchemez S, Reed C, Gambhir M, Finelli L (2014) Estimates of the reproduction number for seasonal, pandemic, and zoonotic influenza: a systematic review of the literature. *BMC Infect Dis* 14(1):480. <https://doi.org/10.1186/1471-2334-14-480>
- Britton T, Becker NG (2000) Estimating the immunity coverage required to prevent epidemics in a community of households. *Biostatistics* 1(4):389–402
- Centers for Disease Control: Seasonal influenza vaccine safety: a summary for clinicians (2016). [https://www.cdc.gov/flu/professionals/vaccination/vaccine\\_safety.htm](https://www.cdc.gov/flu/professionals/vaccination/vaccine_safety.htm)
- Centers for disease control: CDC reports this season's flu vaccine reducing risk by nearly half (2017). <https://www.cdc.gov/flu/spotlights/flu-vaccine-reducing-risk.htm>
- Chao DL, Halloran ME, Obenchain VJ, Longini IM Jr (2010) Flute, a publicly available stochastic influenza epidemic simulation model. *PLOS Comput Biol* 6(1):1–8. <https://doi.org/10.1371/journal.pcbi.1000656>
- Chapman GB, Coups EJ (1999) Predictors of influenza vaccine acceptance among healthy adults. *Prevent Med* 29(4):249–262
- Ferguson NM, Cummings DAT, Fraser C, Cajka JC, Cooley PC, Burke DS (2006) Strategies for mitigating an influenza pandemic. *Nature* 442(7101):448–452. <https://doi.org/10.1038/nature04795>
- Fine PE (1993) Herd immunity: history, theory, practice. *Epidemiol Rev* 15(2):265–302
- Geoffard PY, Philipson T (1997) Disease eradication: private versus public vaccination. *Am Econ Rev* 87(1):222–230. <https://doi.org/10.2307/2950864>
- Hardin R (2013) The free rider problem. In: Zalta EN (ed) *The Stanford encyclopedia of philosophy*, spring, 2013th edn. Stanford University, Metaphysics Research Lab, Cambridge
- Iorio AM, Bistoni O, Galdiero M, Lepri E, Camilloni B, Russano AM, Neri M, Basileo M, Spinozzi F (2012) Influenza viruses and cross-reactivity in healthy adults: humoral and cellular immunity induced by seasonal 2007/2008 influenza vaccination against vaccine antigens and 2009 A(H1N1) pandemic influenza virus. *Vaccine* 30(9):1617–1623. <https://doi.org/10.1016/j.vaccine.2011.12.107>
- Kermack WO, McKendrick AG (1927) A contribution to the mathematical theory of epidemics. *Proc R Soc Lond Ser A* 115(772):700–721
- Krammer F (2019) The human antibody response to influenza A virus infection and vaccination. *Nat Rev Immunol* 19(6):383–397
- Lee J, Paparoditis P, Horton AP, Früwirth A, McDaniel JR, Jung J, Boutz DR, Hussein DA, Tanno Y, Pappas L, Ippolito GC, Corti D, Lanzavecchia A, Georgiou G (2019) Persistent antibody clonotypes dominate the serum response to influenza over multiple years and repeated vaccinations. *Cell Host Microbe* 25(3):367–376.e5. <https://doi.org/10.1016/j.chom.2019.01.010>
- Li M, Taylor EG, Atkins KE, Chapman GB, Galvani AP (2016) Stimulating influenza vaccination via prosocial motives. *PLoS ONE* 11(7):e0159-780
- Loomba S, de Figueiredo A, Piatek SJ, de Graaf K, Larson HJ (2021) Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nat Hum Behav* 5(3):337–348
- Ma J, Earn DJD (2006) Generality of the final size formula for an epidemic of a newly invading infectious disease. *Bull Math Biol* 68(3):679–702. <https://doi.org/10.1007/s11538-005-9047-7>
- Moa AM, Chughtai AA, Muscatello DJ, Turner RM, MacIntyre CR (2016) Immunogenicity and safety of inactivated quadrivalent influenza vaccine in adults: a systematic review and meta-analysis of randomised controlled trials. *Vaccine* 34(35):4092–4102. <https://doi.org/10.1016/j.vaccine.2016.06.064>
- Molinari NAM, Ortega-Sanchez IR, Messonnier ML, Thompson WW, Wortley PM, Weintraub E, Bridges CB (2007) The annual impact of seasonal influenza in the us: measuring disease burden and costs. *Vaccine* 25(27):5086–5096. <https://doi.org/10.1016/j.vaccine.2007.03.046>



- Oraby T, Bauch CT (2015) Bounded rationality alters the dynamics of paediatric immunization acceptance. *Sci Rep* 5(10):724
- Oraby T, Thampi V, Bauch CT (2014) The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proc R Soc B Biol Sci* 281(1780):20133–172
- Osterholm MT, Kelley NS, Sommer A, Belongia EA (2012) Efficacy and effectiveness of influenza vaccines: a systematic review and meta-analysis. *Lancet Infect Dis* 12(1):36–44. [https://doi.org/10.1016/S1473-3099\(11\)70295-X](https://doi.org/10.1016/S1473-3099(11)70295-X)
- Parker AM, Vardavas R, Marcum CS, Gidengil CA (2013) Conscious consideration of herd immunity in influenza vaccination decisions. *Am J Prevent Med* 45(1):118–121
- Prosper O, Saucedo O, Thompson D, Torres-Garcia G, Wang X, Castillo-Chavez C (2011) Modeling control strategies for concurrent epidemics of seasonal and pandemic H1N1 influenza. *Math Biosci Eng* 8(1):141–170
- Rambhia KJ, Rambhia MT (2018) Early bird gets the flu: what should be done about waning intraseasonal immunity against seasonal influenza? *Clin Infect Dis* 68(7):1235–1240. <https://doi.org/10.1093/cid/ciy748>
- Taubenberger JK, Kash JC (2010) Influenza virus evolution, host adaptation, and pandemic formation. *Cell Host Microbe* 7(6):440–451
- Thaler RH, Sunstein CR (2009) *Nudge: improving decisions about health, wealth, and happiness*. Penguin, New York
- Tversky A, Kahneman D (1974) Judgement under uncertainty: heuristics and biases. *Science* (80- ) 185(4157):1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Verelst F, Willem L, Beutels P (2016) Behavioural change models for infectious disease transmission: a systematic review (2010–2015). *J R Soc Interface* 13(125):20160–820
- Voinson M, Billiard S, Alvergne A (2015) Beyond rational decision-making: modelling the influence of cognitive biases on the dynamics of vaccination coverage. *PLoS ONE* 10(11):e0142-990
- Wells CR, Bauch CT (2012) The impact of personal experiences with infection and vaccination on behaviour-incidence dynamics of seasonal influenza. *Epidemics* 4(3):139–151
- World Health Organization: Questions and answers—vaccine effectiveness estimates for seasonal influenza vaccines (2015). [http://www.who.int/influenza/vaccines/virus/recommendations/201502\\_qanda\\_vaccineeffectiveness.pdf](http://www.who.int/influenza/vaccines/virus/recommendations/201502_qanda_vaccineeffectiveness.pdf)
- World Health Organization: Up to 650 000 people die of respiratory diseases linked to seasonal flu each year (2017). <https://www.who.int/news-room/detail/14-12-2017-up-to-650-000-people-die-of-respiratory-diseases-linked-to-seasonal-flu-each-year>

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