Supplemental information

Interfering with Influenza: nonlinear coupling of reactive and static mitigation strategies

published by Journal of the Royal Society Interface

Cameron Zachreson, Kristopher M. Fair, Nathan Harding, Mikhail Prokopenko

March 24, 2020

S1 Detailed methods

S1.1 Population demographics

We used the Australian Bureau of Statistics (ABS) 2016 Census to generate a collection of Agents with the demographic characteristics of Australia's population. The demographic and housing data we use in our model is aggregated on the level of Statistical Area Level One (SA1), which are statistical geographical regions containing 200-800 residents (these are referred to as 'neighbourhoods' in the main text). This is one of the primary levels on which the ABS partitions demographic (e.g., age, sex, employment status) and housing (e.g., household size and composition) statistics. In the population generation method used here, these statistics are interpreted as (dependent) probability distributions describing Agent properties according to their region of residence.

The process of stochastically generating Agents for the population begins by iterating through SA1s. Before generating Agents, a household is generated based on the ABS partition statistics for a given SA1. The possible household types are lone, group, or family types. These are then further subdivided into the following categories: single, couples with children, couples without children, or single parent families. Finally, the household is assigned a size conditional on both its type, and the distribution of household sizes in that SA1. At this point, Agents are generated as adults or children depending on the particular households size and type. Additional attributes are assigned using population statistics about age and sex. This process continues until the population of the SA1 is reached, at which point we select another, and continue the process until all SA1 partitions have been populated.

S1.2 Mid- to long-range social contact model

After the Agent population has been generated, workplaces and schools are assigned to adults and children respectively. Workplaces are assigned based on the commuter origin-destination (OD) matrix obtained from the ABS and adjusted according to the procedure outlined in our previous work to correct for some

of the noise introduced by the ABS [1, 2]. Each entry of the OD matrix gives a home SA1 h, a workplace DZN w and a number of individuals who live in h and work in w. As the Agents have already been assigned residential regions, they are chosen from households uniformly at random. The ABS census does not contain data about the location of schools, or where students commute from. Therefore we place schools pseudo-deterministically based on data from the Australian Curriculum, Assessment and Reporting Authority (ACARA) by postal code from a non-census dataset containing the most complete information available on school enrolment numbers and locations since 2008. In the absence of school commuting data, we stochastically assign students to schools based primarily on proximity, assuming that students would attend a school if they live within its catchment zone. If a student is within the catchment zone of multiple schools, the school is chosen randomly based on the number of available places. Further descriptions of the school placement and student assignment methods can be found in our previous works [1, 3].

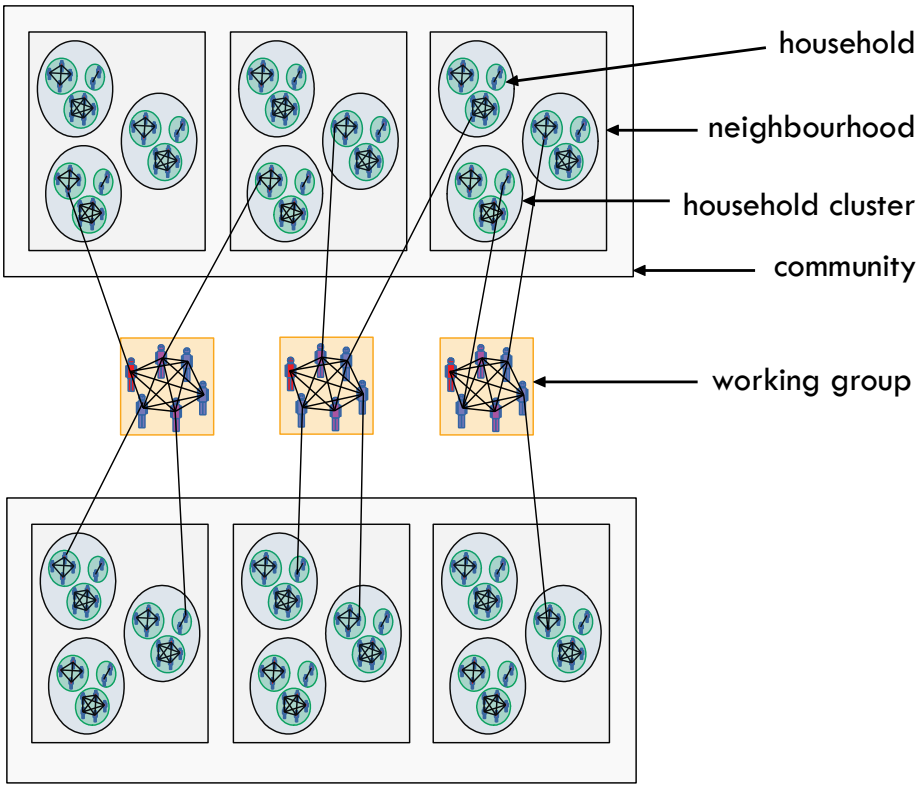

Figure S1: Schematic of the community structure used in the ACEMod simulation platform. Agents mix locally in household, household cluster, neighbourhood, and community groups. Transmission between communities is possible due to interactions in working groups and school groups (which include classroom, grade, and school sub-groups). For simplicity, the school mixing groups are not shown in the schematic. The different colours of the Agents represent the incorporation of demographic factors such as age and sex.

S1.3 Transmission model

In the model of illness and disease transmission used in this study, infected individuals spread the virus to others through several types of mixing groups. These mixing groups describe a regular pattern of behaviour associated with each individual which has the potential to transmit infection, such as regular contact with work colleagues, classmates, and family members. Each simulated day is split into two distinct phases, "daytime" and "night time". During the daytime phase individuals can spread the infection to others in

their working groups which include workplaces for most adults, as well as classes, grades, and schools for teachers and students. During the night time phase, individuals are able to spread infection within a number of mixing groups, such as their household, household cluster, neighborhood and community. For transmission parameters related to each of these mixing groups, see [1].

For each individual i in the population there is a collection of different mixing groups, denoted $\mathcal{G}_i(t)$, where that individual Agent interacts with its contacts at time t. Each mixing group $g \in \mathcal{G}_i(t)$ is associated with a static set of Agents $\mathcal{A}_g$. At an individual level we are interested in the probability of a susceptible Agent i becoming infected during a given time period (step) t. Let $X_i(t)$ denote the state of individual i at time t. The probability of a susceptible Agent i becoming infected during a given time period is given by

$$p_i(t) = Pr\{X_i(t) = \text{LATENT} | X_i(n-1) = \text{SUSCEPTIBLE}\} \tag{S1}$$

where LATENT denotes the first stage of infection. The infection probability of a susceptible individual i is computed as

$$p_i(t) = 1 - \prod_{g \in \mathcal{G}_i(t)} \left[\prod_{j \in \mathcal{A}_g \setminus i} \left(1 - p^g_{j \to i}(t)\right) \right] \tag{S2}$$

where $\mathcal{G}_i(t)$ denotes the mixing groups where Agent i interacts at time step t, and $p^g_{j \to i}$ denotes the instantaneous probability of transmission from Agent j to i in contact group g. An Agent has a set probability of becoming either symptomatic or asymptomatic. For an asymptomatic Agent, the probability of transmission to a susceptible Agent is halved.

The context-dependent probabilities of transmission $p^g_{j \to i}$ are globally scaled by a transmission multiplier κ, which modulates the potential for the epidemic to spread. They are also dependent on the natural history of the disease, which determines the relative infectiousness f of an individual as a function of the time elapsed since infection:

$$p^g_{j \to i}(t) = \kappa f(t - t_j | j, i) q^g_{j \to i}\,, \tag{S3}$$

where context-dependence is given by the $q^g_{j \to i}$ terms, which were calibrated from empirical studies where possible. However, in the absence of known context-dependent transmission probabilities, additional steps were taken to calculate them from known contact rates. For additional details of this process see our existing work [1].

S1.4 Mitigation strategies

The likelihood of an individual Agent becoming infected each day is dependent on their contact network for the daytime/night-time phase and the transmission probability produced by infected Agents within that network. In practice, intervention strategies influence the contagion dynamics by effecting either the transmission probability or contact network to suppress the outbreak.

S1.4.1 Vaccination

The vaccination protocol is applied uniformly at random to each Agent based on a fixed probability corresponding to the compliance rate specified at the start of each simulation. If a vaccinated Agent is infected, we alter the contribution to the force of infection on their neighbors $p^g_{i \to j}(t)$ (the infectiousness). For susceptible, vaccinated Agents, we increase their resilience to transmission $p_i(n)$ (the susceptibility). The vaccine efficacy for susceptibility, VE_S, determines a multiplier θ for the infected Agent's contribution to the force of infection related by $\theta = 1 - VE_S$. Similarly, an uninfected Agent's probability of infection

is altered by the multiplier η, determined by the vaccine efficacy of infectiousness $\eta = 1 - VE_I$. We have made the assumption of a low-efficacy vaccination, as used in Longini *et al.*, with $VE_S = 0.30$ and $VE_I = 0.50$ [4].

S1.4.2 Antivirals

The antiviral prophylaxis protocol functions in a similar manner to vaccination, with a corresponding antiviral efficacy of infectiousness AVE_I and antiviral efficacy of susceptibility AVE_S that reduce the infectiousness and susceptibility (respectively) of Agents while dosed with antivirals. In addition to this, antivirals reduce the illness recovery time from 5.5 days to 4.5 days, and reduce the symptomatic fraction. The latter effect introduces the antiviral efficacy for illness $AVE_D = 1 - \Phi$. Here Φ (the probability of an Agent remaining asymptomatic when infected) is increased above its usual value of 0.331, due to the antivirals. We use antiviral efficacy of the neuraminidase inhibitor Oseltamivir, $AVE_S = 0.30$, $AVE_I = 0.62$ and $AVE_D = 0.60$ [5, 6].

Unlike vaccination, the application of AVP is done dynamically after a symptomatic infection is detected. Each day an undetected symptomatic Agent has a probability of being detected ($p_{detect} = 0.1$). If detected, the Agent (index case) will be given antivirals after 1 day and Agents that are considered at risk of infection from the index case are given antivirals after the designated response time has elapsed. This delayed response represents a logistical roll out time required to implement localized pandemic intervention programs. Here, we consider two tactics for identifying at-risk contacts of index cases. The first is contact-targeted antiviral prophylaxis (TAP), in which all members of the infected Agent's household and work group are given antiviral prophylaxis (for school-aged children and teachers, the work group targeted is their entire school). The second is geographically targeted antiviral prophylaxis (GTAP), in which all Agents residing in the same SA1 as the index case receive treatment (these correspond to the neighborhoods depicted in Figure S1).

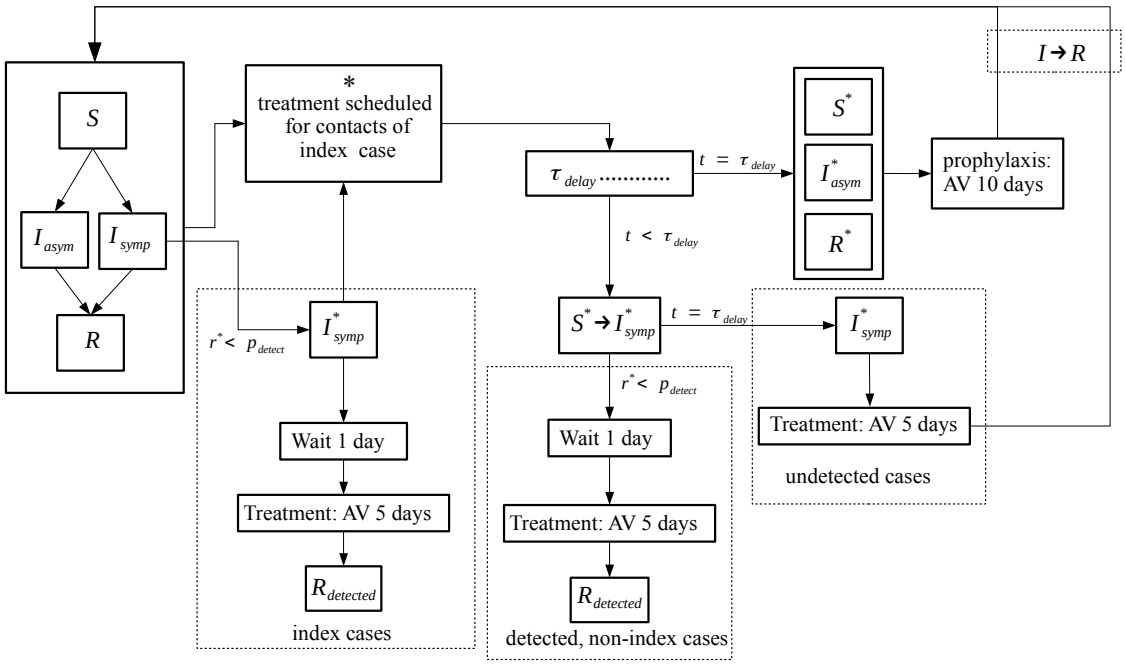

Figure S2: Schematic workflow of intervention implementation. The associated symbols and their meanings are given in Table S1. If index cases are detected, that event triggers prophylaxis of their contacts or neighbors after a delay time associated with logistical implementation of the measures. Note that disease detection and treatment can occur independently of the large-scale intervention protocol if illness is detected before the delay time is up.

Table S1: Parameters and states associated with the integrated SIR transmission, illness detection, and intervention models.

Symbol	description
S	Agent state - Susceptible (can become infected)
I_{asym}	Agent state - Infected, asymptomatic (cannot be detected)
I_{symp}	Agent state - Infected, symptomatic (may be detected)
R	Agent state - Recovered (undetected - may be scheduled for prophylaxis)
I^*_{symp}	Agent state - Infected, scheduled for treatment (treated after 1 day if detected, treated after delay period if undetected)
$R_{detected}$	Agent state - Recovered, detected (will not be scheduled for prophylaxis)
S^*	Agent state - Susceptible, scheduled for prophylaxis after delay period
I^*_{asym}	Agent state - Infected, asymptomatic, scheduled for prophylaxis after delay period
R^*	Agent state - Recovered, scheduled for prophylaxis after delay period
r^*	pseudorandom number ($r^* \in [0, 1]$)
p_{detect}	probability that a symptomatic infection will be detected ($p_{detect} \in [0, 1]$)
τ_{delay}	delay period (in days) between index case detection and targeted intervention
t	time (in days) since detection

S2 Supplementary results

S2.1 Reproductive Ratio R_0

Here we describe our procedure for estimating the reproductive ratio as a function of the global infection scalar κ. We compare the method used in this work (the "Attack Rate Pattern Weighted" method) with an alternative uniform sampling method. Briefly, both methods estimate the reproductive ratio by randomly sampling individuals from the population and counting the number of secondary infections resulting from that individual becoming infected. By repeating such a process several thousand times, an estimate of the reproductive ratio is possible.

Due to population heterogeneity, biases are introduced if this sampling is performed uniformly at random. This is because different cohorts of the population are more or less likely to become infected based on the population structure and differences in susceptibility and infectiousness. Heterogeneity is particularly prominent in the clusters of individuals from the 5 to 18 year-old age group found in schools. Due to the strong correlation between age group and population structure, it is reasonable to expect that biasing the sampling of index cases based on the age-stratified contribution to the overall attack rate would give an improved estimate of the reproductive ratio.

The plots in Figure S3 confirm this expectation. Figure S3(a) demonstrates the discrepancy between the value computed by the uniform random sampling method R_{rand} and the alternative, weighted estimate R_{ARPW}, as functions of the global infection scalar κ. The reason for the discrepancy is made clear by examining Figure S3(b), which shows the non-uniform contributions to the overall attack rate, stratified by age-group, as functions of κ. Lastly, Figures S3(c) and S3(d) plot the attack rate as functions of R_{rand} and R_{ARPW}, respectively. These plots demonstrate how the Attack Rate Pattern-Weighted method produces a more intuitive result that is consistent with the conceptual definition of the reproductive ratio as the expected number of secondary cases produced by a typical index case. For the data reported in Figure S3(a) we averaged R_{ARPW} from $(5 \times 10^3 < n < 1.5 \times 10^4)$ and R_{rand} from $(5 \times 10^4 < n < 6 \times 10^4)$ index cases for each value of κ.

S2.2 Pre-pandemic vaccination

To demonstrate that our vaccination programs perform as expected for various values of R_0 and compliance, Figure S4(a) shows that employing uniform pre-pandemic vaccination suppresses the reproductive speed of the virus, both delaying and decreasing the peak incidence of new illnesses (symptomatic infections). On the same point, Figure S4(b) demonstrates how this effect extends to the threshold behavior of epidemic dynamics: with uniform vaccination, the R_0 threshold for extensive epidemics moves to higher values.

Note that in the absence of vaccination, the critical value of R_0 is less than 1 (which is the critical value in mean-field models by definition). This departure from continuum dynamics occurs for two reasons: firstly, because of the heterogeneous contact structure in our population model [7], and secondly, because index cases are produced continuously through arrivals from overseas which guarantees finite attack rates. Because of these factors, there is a non-vanishing probability that an individual with many more contacts than a typical member of the population will become infected and transmit to two or more of their immediate contacts, who due to social clustering (i.e., school and classroom groups) are also likely to have higher than average contact numbers.

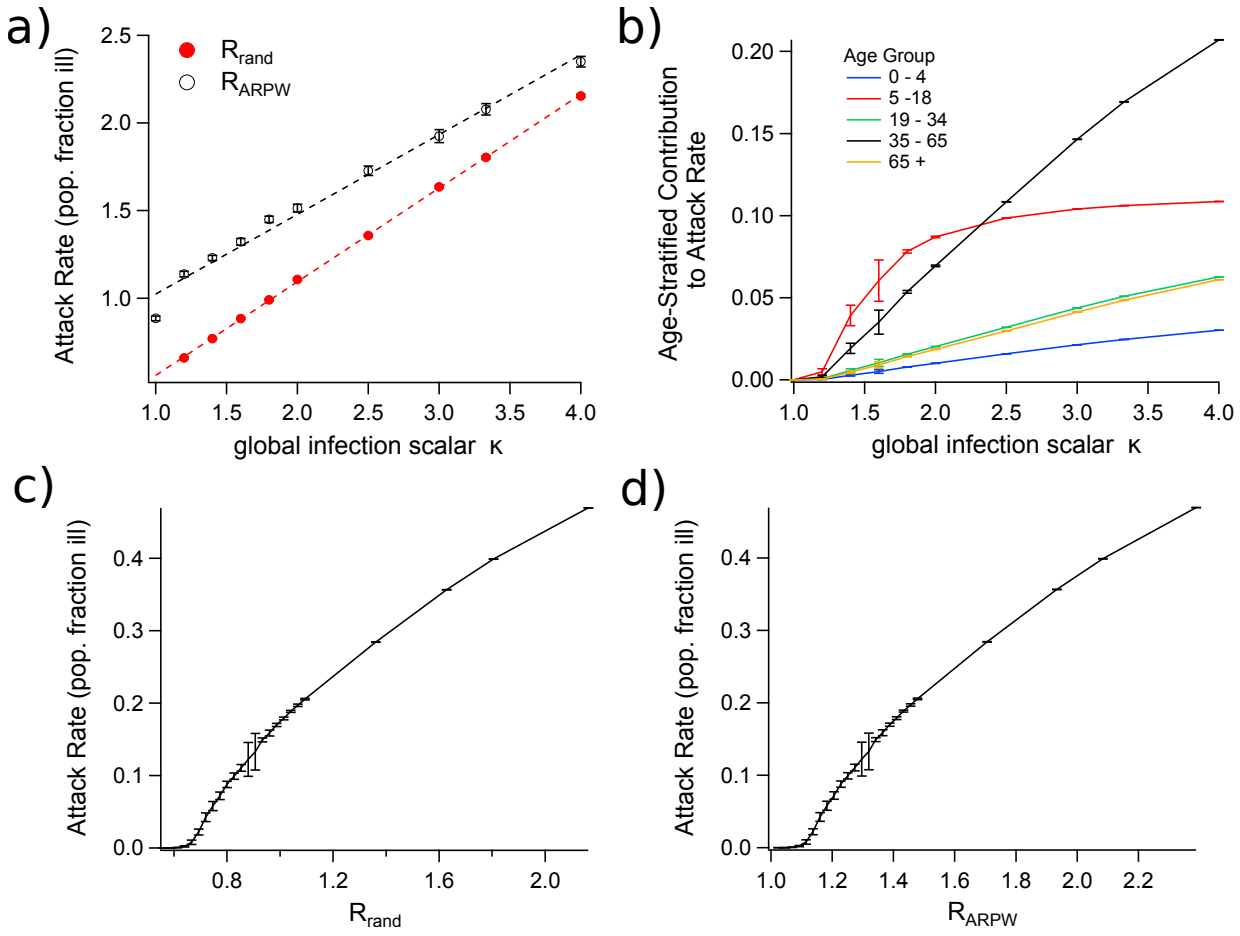

Figure S3: Comparison of reproductive ratio estimates based on uniform random sampling (R_{rand}) and attack rate pattern weighted (R_{ARPW}) methods. (a) Plots of the two estimates as functions of the infectiousness parameter κ. Black and red dashed lines are linear fits for R_{ARPW} and R_{rand} as functions of κ ($R^2 = 0.9848$ and $R^2 = 0.9998$, respectively; error bars: $\pm$ SEM). (b) Age-stratified contributions to the attack rate as functions of κ. (c) Attack rate as a function of R_{rand}. (d) Attack rate as a function of R_{ARPW}. [Error bars in (b), (c), and (d) show $\pm$ std., R_{rand} and R_{ARPW} values in (c) and (d) were determined from κ using the linear fits shown in (a).]

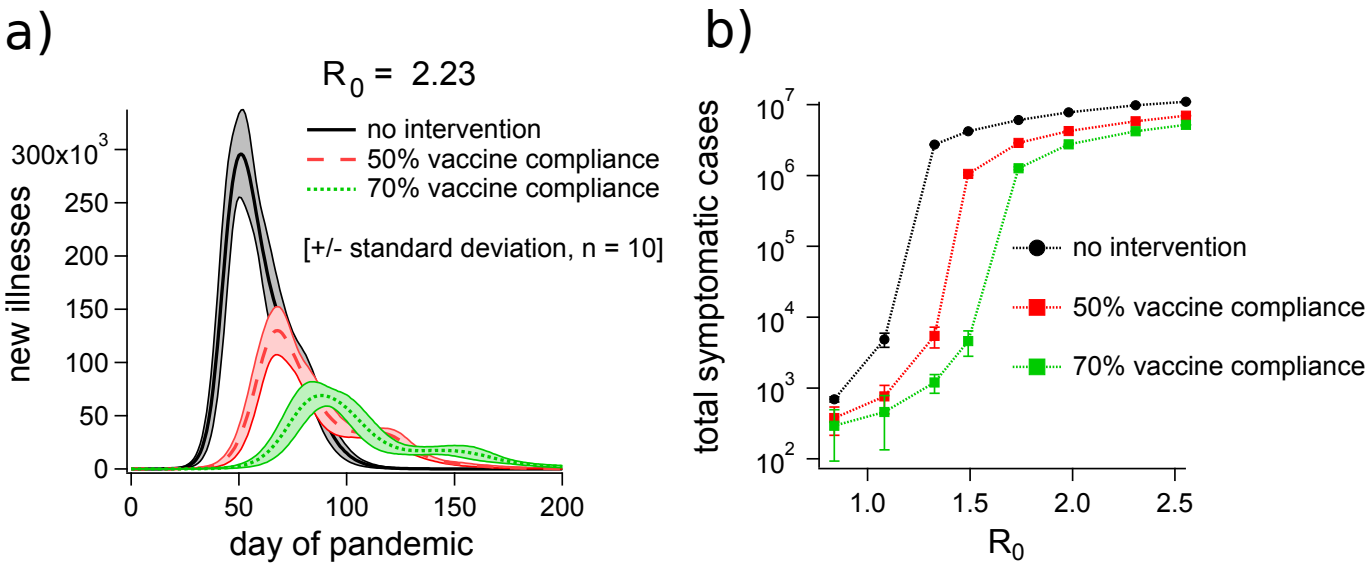

Figure S4: Performance of simulated pre-pandemic vaccination programs, for different values of compliance and disease infectiousness. (a) Epidemic curves (symptomatic cases vs. time) for vaccination compliance rates of 0% (gray), 50% (red), and 70% (green) with fixed reproductive ratio $R_0 = 2.23$ [lines represent mean values, colored bands show $\pm$ standard deviations over 10 runs]. (b) Plots of attack rate (total illnesses) as functions of R_0 for vaccination compliance rates of 0% (black circles), 50% (red squares), and 70% (green triangles) [computed values are shown as symbols, the lines connecting them are guides to the eye, error bars show $\pm$ standard deviation, n $= 10$].

S2.3 Geographically-targeted antiviral prophylaxis

Contact targeted prophylaxis carries with it the unfortunate requirement of sufficiently complete contact tracing that antivirals can be distributed to the social contacts of index cases. The feasibility of such strategies has been questioned, and other targeting approaches suggested [4]. One of these is to simply administer antivirals to the residential neighborhoods of index cases, which, in some communities and cultures could effectively account for the majority of index case contacts. However, as our results demonstrate, such targeting methods are not effective in large-scale urban environments in which long-distance commuting is the norm. Australia has a very high urban population fraction, and a very centralized urban economic structure, with large sectors of the workforce commuting to small areas of central business districts from disparate suburbs on a daily basis. It is perhaps unsurprising then, that GTAP is substantially less efficient than TAP, in this context, as our results demonstrate.

Figure S5 presents a report of GTAP behavior as a function of delay time and compliance rate. The first clear difference from the TAP protocol is that much larger administration numbers are required for implementation of the program. However, given a very large stockpile, the results demonstrate that higher suppression levels are possible. This is intuitive: given that neighborhoods are typically larger than social clusters, the GTAP protocol allows for a more universal coverage with antivirals, however, the program is prohibitively inefficient in the areas of parameter space for which suppression is substantial.

In addition to this general finding, there are several other specific points that can be gleaned from the results:

- The antiviral administration dynamics are multi-modal. Comparison of Figures S5(a), (b), and (c) shows that the interval between successive maxima in antiviral administration numbers lengthens with the logistical delay time.
- Despite higher suppression levels, the effect on the growth rate of the virus and incidence peak timing is still sensitive to delay time.
- As for the TAP program, Figure S5(e) shows that the efficiency of GTAP is non-monotonic with compliance for short delay times, and generally decreases with higher compliance rates.

These three points have a coherent implication: the negative effect of local reductions in transmission rate on mitigation efficiency are more pronounced when antivirals are distributed geographically. Because GTAP does not account for long-range transmission through interactions outside of residential communities, it has weaker effects on the geographic spread of the disease. However, its large-scale local effects slow the virus substantially during the prophylaxis period (particularly when the delay time is small). These two factors combine to strongly promote the local resurgence of the virus in the same communities, leading to many rounds of prophylaxis and requiring very large outlay of resources to suppress illness levels.

Taken together, these results indicate that GTAP by itself (as implemented in our simulations) is not a feasible strategy for pandemic mitigation, as it would require prohibitively large antiviral stockpiles, and the population would be subject to many rounds of prophylaxis that would severely impact quality of life. Of course, under some circumstances such a scenario may still be preferable, depending on case severity, and the feasibility of the contact tracing required for an effective TAP protocol.

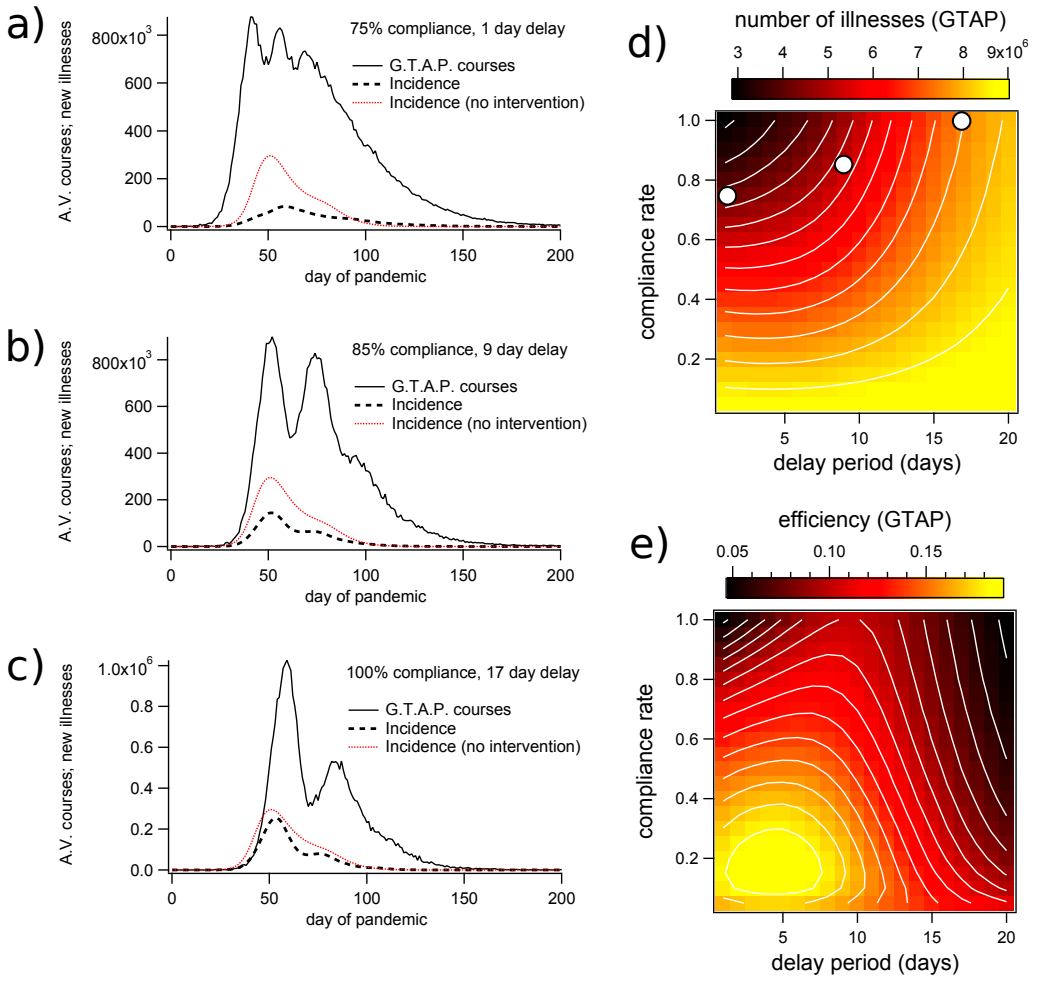

Figure S5: Simulation results for geographically-targeted antiviral prophylaxis (GTAP). The simulation parameters used to generate this data are identical to those presented in Figure **??**, the only difference is the type of targeting strategy employed. Plots (a - c) show the results of single simulation instances, with reference plots of illness incidence when no intervention is applied. Plots (d) and (e) show attack rates and efficiencies (respectively) as functions of compliance with the GTAP protocol within targeted localities.

S2.3.1 Geographically-targeted antiviral prophylaxis with pre-pandemic vaccination

Unfortunately, the beneficial effects observed from combining TAP with pre-pandemic vaccination do not apply to GTAP as well. The plots of illness incidence and GTAP courses in time shown in Figures S6(a - c) illustrate that while attack rate suppression is still substantial, the effect on the initial epidemic growth rate and peak timing are not. Additionally, the number of antiviral courses deployed is till extremely high, despite the greatly suppressed attack rate. In fact, comparison of Figure **??**(d) and Figure S6(d) shows that, when combined with vaccination, GTAP performance is inferior to that of TAP, even when taking into account only total suppression levels. The situation deteriorates further when efficiency is taken into account. As shown in Figure S6(e), there is no reversal of the diminishing return on GTAP compliance, and there is no local maximum of efficiency for intermediate compliance rates. All in all, the argument for the GTAP program is not strong in any area of the parameter space explored here. While it is capable of dramatic suppression of illness levels, it does this at great expense, and shows no promise of beneficial coupling with the pre-pandemic vaccine program.

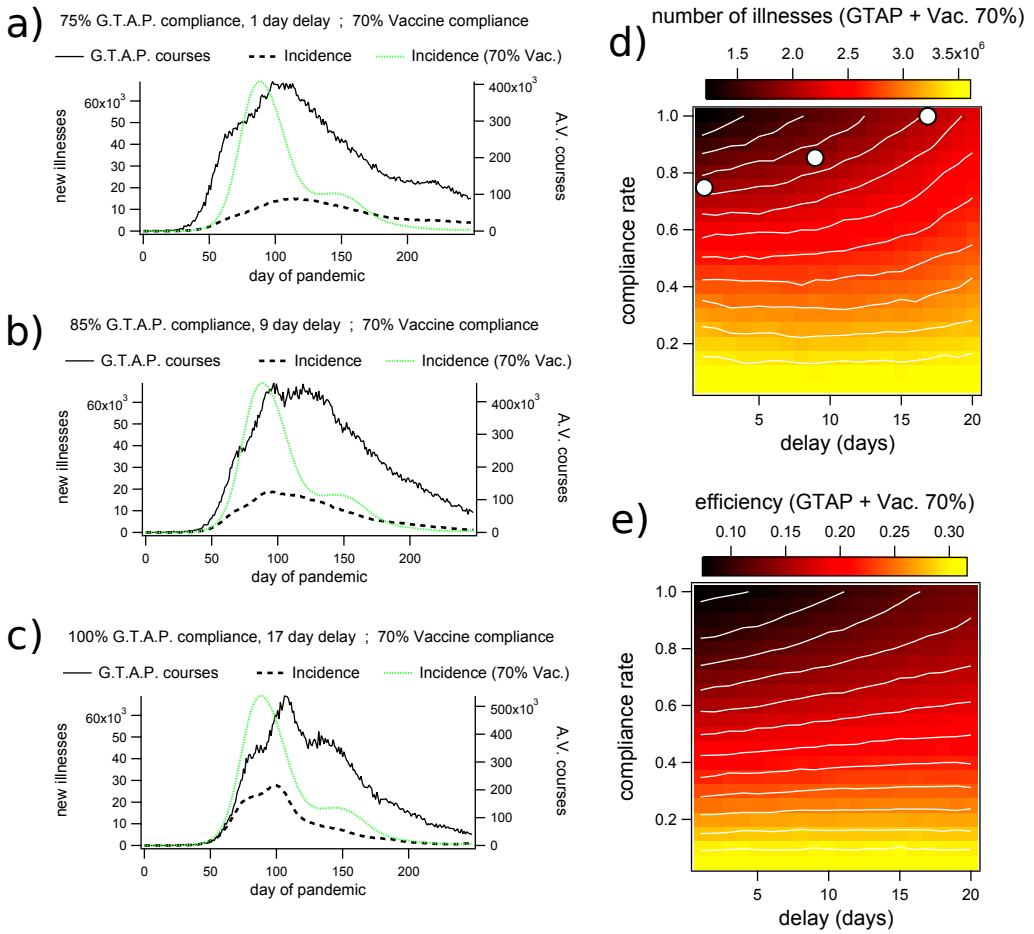

Figure S6: Simulation results for geographically targeted antiviral prophylaxis (GTAP) in combination with 70% pre-pandemic vaccination. Plots (a - c) show results for single simulation instances using different values of GTAP compliance and delay time, with illness incidence computed for 70% vaccine compliance plotted in green for reference. Note that antiviral courses as a function of time are plotted to the right axis, and are approximately one order of magnitude larger than the corresponding illness incidence. Plots (d) and (e) plot attack rate and net suppression per course (summing antiviral courses and vaccines deployed) as functions of delay period and GTAP compliance in targeted localities.

S2.4 Antiviral administration numbers

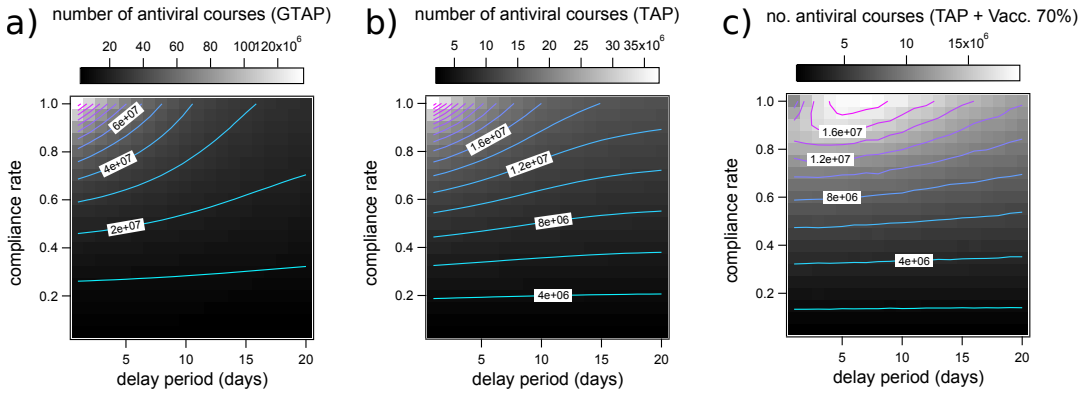

Figure S7: Two-dimensional image plots of the total number of doses administered for three dynamic mitigation protocols: geographically targeted antiviral prophylaxis (a), contact targeted antiviral prophylaxis (b), and contact targeted antiviral prophylaxis with 70% of the population vaccinated prior to pandemic onset (c), as functions of compliance rate and delay time. Note that (c) reports the number of antiviral courses only.

S2.5 peak timing

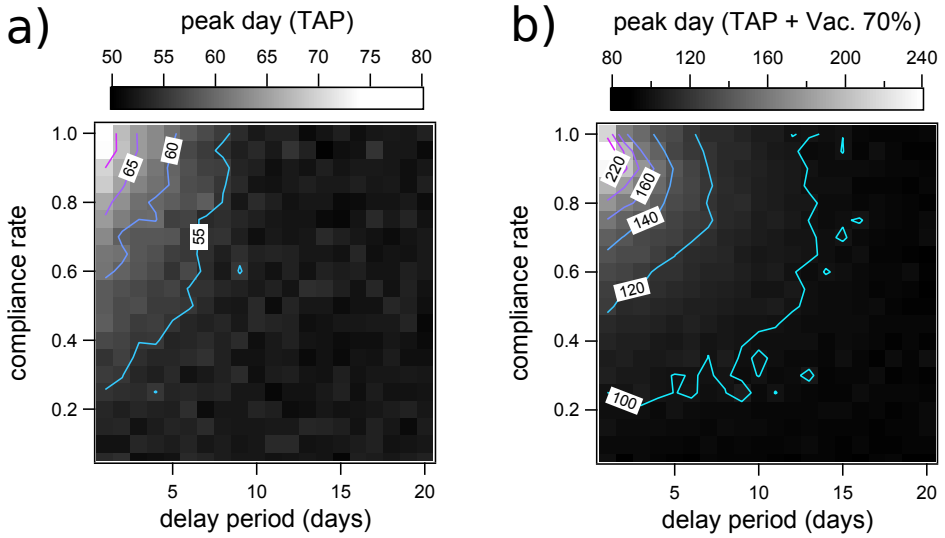

Figure S8: Two-dimensional image plots of peak timing (the day corresponding to the maximum number of concurrent symptomatic cases) for contact-targeted antiviral prophylaxis (a) and contact-targeted antiviral prophylaxis with 70% of the population vaccinated prior to pandemic onset (b), as functions of compliance rate and delay time.

S3 Supplementary movies

These movies show the spatiotemporal spread of the pandemic from single simulation instances over maps of Australia, with details of the regions surrounding Melbourne (Vic) and Sydney (NSW). The colors correspond to normalized prevalence values for each statistical region on the level of Statistical Area Level 2 (SA2), which is one step up in the spatial hierarchy from the SA1 level on which the population statistics are sampled in our population model. For display purposes we normalized the prevalence data to the maximum local prevalence value computed for the corresponding simulation, and the colors correspond to the natural logarithm of these normalised prevalence levels. In addition to the maps, each movie contains the corresponding plot of aggregated national prevalence and antiviral deployments in time. All movies correspond to $R_0 = 2.23$, and each has a different set of mitigation procedures and related parameters. For reference, a choropleth of the Australian population distribution is also included in Figure S9.

- Supplemental Movie SM1 (SM1_No_Intervention.pdf):
  This movie represents the baseline condition, where no mitigation efforts are employed.
- Supplemental Movie SM2 (SM2_Vaccination_70.pdf):
  This movie corresponds to 70% of the population vaccinated prior to the pandemic onset.
- Supplemental Movie SM3 (SM3_TAP75_1day_delay.pdf):
  This movie corresponds to dynamic, contact-targeted antiviral prophylaxis with a compliance rate of 75% and a delay time of one day.
- Supplemental Movie SM4 (SM4_GTAP75_1day_delay.pdf):
  This movie corresponds to dynamic, geographically targeted antiviral prophylaxis with a compliance rate of 75% and a delay time of one day.
- Supplemental Movie SM5 (SM5_PVac70_TAP75_1day_delay.pdf):
  This movie corresponds to dynamic, contact-targeted antiviral prophylaxis with a compliance rate of 75% and a delay time of one day. Prior to pandemic onset, the population is vaccinated at a compliance rate of 70%
- Supplemental Movie SM6 (SM6_PVac70_TAP85_9day_delay.pdf):
  This movie corresponds to dynamic, contact-targeted antiviral prophylaxis with a compliance rate of 85% and a delay time of nine days. Prior to pandemic onset, the population is vaccinated at a compliance rate of 70%
- Supplemental Movie SM7 (SM7_PVac70_GTAP75_1day_delay.pdf):
  This movie corresponds to dynamic, geographically targeted antiviral prophylaxis with a compliance rate of 75% and a delay time of one day. Prior to pandemic onset, the population is vaccinated at a compliance rate of 70%
- Supplemental Movie SM8 (SM8_PVac70_GTAP85_9day_delay.pdf):
  This movie corresponds to dynamic, geographically targeted antiviral prophylaxis with a compliance rate of 85% and a delay time of nine days. Prior to pandemic onset, the population is vaccinated at a compliance rate of 70%

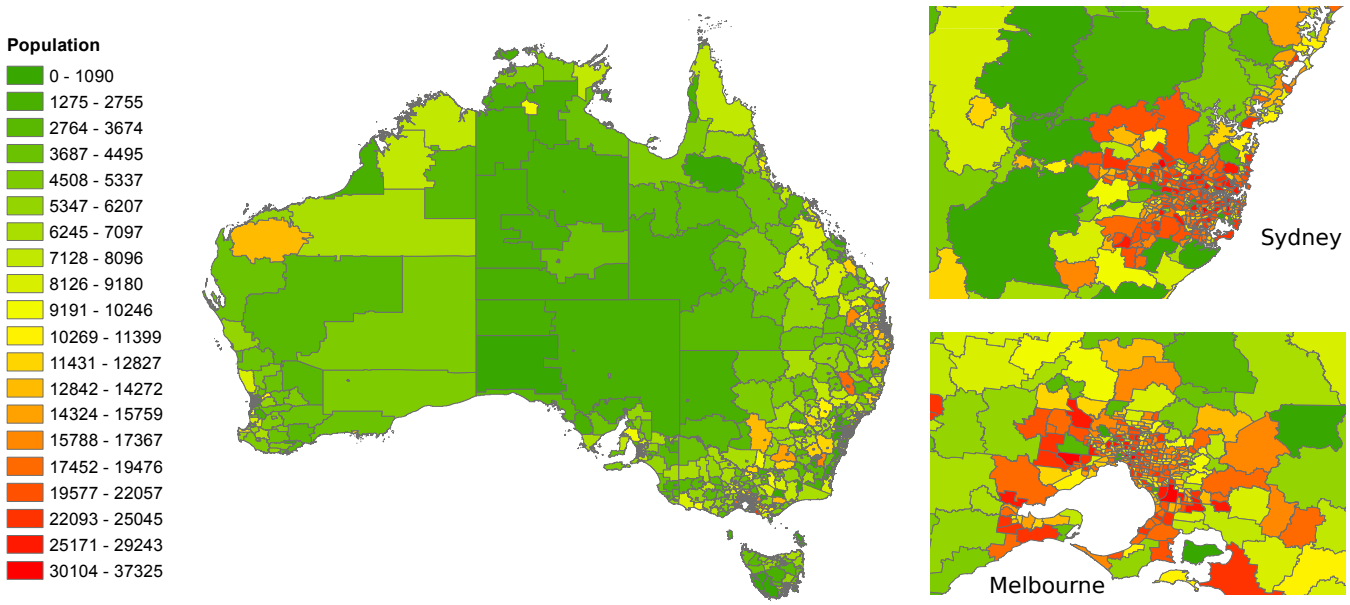

Figure S9: **Australian population distribution (2016):** A choropleth illustrating the heterogeneous population distribution of Australia from the year 2016, at the scale of Statistical Area Level 2 (SA2). This is provided as a reference for interpreting the data displayed in the Supplemental Movies.

References

[1] Cliff OM, Harding N, Piraveenan M, Erten EY, Gambhir M, Prokopenko M. Investigating spatiotemporal dynamics and synchrony of influenza epidemics in Australia: An agent-based modelling approach. Simul Model Pract Th. 2018;87:412–431.

[2] Fair KM, Zachreson C, Prokopenko M. Creating a surrogate commuter network from Australian Bureau of Statistics census data. Scientific data. 2019;6(1):1–14.

[3] Zachreson C, Fair KM, Cliff OM, Harding N, Piraveenan M, Prokopenko M. Urbanization affects peak timing, prevalence, and bimodality of influenza pandemics in Australia: Results of a census-calibrated model. Sci Adv. 2018;4(12):eaau5294.

[4] Longini IM, Nizam A, Xu S, Ungchusak K, Hanshaoworakul W, Cummings DA, et al. Containing pandemic influenza at the source. Science. 2005;309(5737):1083–1087.

[5] Welliver R, Monto AS, Carewicz O, Schatteman E, Hassman M, Hedrick J, et al. Effectiveness of oseltamivir in preventing influenza in household contacts: a randomized controlled trial. Jama. 2001;285(6):748–754.

[6] Hayden FG, Gubareva LV, Monto AS, Klein TC, Elliott MJ, Hammond JM, et al. Inhaled zanamivir for the prevention of influenza in families. New England Journal of Medicine. 2000;343(18):1282–1289.

[7] Liu QH, Ajelli M, Aleta A, Merler S, Moreno Y, Vespignani A. Measurability of the epidemic reproduction number in data-driven contact networks. Proceedings of the National Academy of Sciences. 2018;115(50):12680–12685.