Supporting Text

A. Simulation Model Details

The three basic elements of our national-level simulation model are (*i*) a previously developed stochastic agent-based model for disease spread at the community level; (*ii*) detailed U.S. Census demographics and worker flow data for daily commuter traffic at short distances and Bureau of Transportation Statistics data for less frequent long-range travel behavior; and (*iii*) high-performance parallel computing expertise in modeling millions to billions of particles on hundreds to thousands of processors. These three components, each of which we describe below in some detail, are brought together to provide a unique capability for a detailed modeling of disease spread in the U.S. population.

1. Community-Level Stochastic Simulation Model

Population structure

As the starting point for constructing our national simulation model, we use a discrete-time stochastic simulation model of disease spread within a structured 2,000-person community. Similar models have been developed and applied previously to both influenza (1-4) and smallpox (5). The model population is stochastically generated to match census-based nationwide distributions of age, household size, and employment status. Each person in the population belongs to one of five age groups: preschool-age children (0-4 years), school-age children (5-18 years), young adults (19-29 years), adults (30-64 years), and older adults (64+ years). Households consist of one to seven persons, with either one or two adults, and are grouped randomly into clusters of four households each, and further grouped into one of four nonoverlapping neighborhoods, each containing ≈500 people. Every person also belongs to a set of close and casual contact (also referred to as “mixing”) groups, ranging from their household and household cluster (highest contact rates) to schools and workplaces, down to their neighborhood and the

entire community (with the lowest contact rates, representing occasional interactions in malls, supermarkets, and churches, for instance).

All preschool-age children are assigned to either a neighborhood daycare center, with 14 children on average, or to one of several smaller neighborhood playgroups, each with four children. Depending on their age, school-age children may belong to one of two elementary school groups (each shared between two neighborhoods, with 79 students each on average), to a community-wide middle school group (128 students on average), or to a community-wide high school group (average 155 students). These school contact groups are in general not actual schools but rather representative of the typical daily interactions a student may have with classmates and other peers. According to U.S. Census data, 93% of children 5–18 years old attend school, so we allow the remaining 7% to mix in the household, household cluster, neighborhood, and community during the daytime. Working adults (restricted to those who are 19–64 years old) belong to a work group of ≈20 people. Although in reality many work places are larger than 20 people, we assume that workers make a contact of sufficient duration and/or closeness to transmit influenza virus with a subset of the entire workforce at that location.

Disease transmission model

Transmission within each contact group is described by a contact probability c_i (Table 3), which may depend on the age of both the infectious and susceptible persons. This contact probability represents the likelihood (within each 12-hour period) of having a contact of sufficient duration and closeness for transmission of an infectious dose of influenza virus to be possible between these two individuals in this social setting. The probability of transmission given such contact, P_{trans}, is a single scalar number that multiplies each contact probability, allowing for a simple variation in contagiousness (typically represented by the basic reproductive number, R_0) without modifying the underlying social interaction network parameters. We do not allow for any seasonal or weekly variation in contact rates or transmission probability, and no births or nonflu-related deaths are included in our model.

Each day, the probability of infection for each susceptible individual is computed based on the transmission probabilities for each potential infectious contact, $p_i = P_{\text{trans}} \times c_i$. If the infectious contact is receiving antiviral treatment, this transmission probability is further multiplied by $(1 - \text{AVE}_i)$, where AVE_i is the antiviral efficacy for infectiousness. Similarly, if they have been vaccinated, the vaccine efficacy for infectiousness VE_i reduces the transmission probability by $(1 - \text{VE}_i)$. The transmission probability p_i can be further reduced for asymptomatic (yet infectious) contacts, as described in the next section. The probability of a susceptible person becoming infected is then computed as a product of all of the possible infectious contacts each day. Fig. 3 illustrates this calculation for a susceptible adult (shown in blue), with one infectious child in the household (HH), one infectious workgroup (WG) contact, and three other infectious people in the wider community (Comm). The probability that this susceptible adult becomes infected is

$$P = 1-(1- p_{HH}^{c\to a})\cdot(1- p_{WG}^{a\to a})\cdot(1- p_{Comm}^{a\to a})\cdot(1- p_{Comm}^{c\to a})^2,$$

where *c* and *a* denote child and adult, respectively. (It is obviously trivial to generalize this to the case where the two community children have different levels of infectiousness due to therapeutic drugs or differing levels of severity, e.g., one symptomatic and the other subclinical.) A Bernoulli trial is conducted by generating a uniform [0,1] random number; if this number is less than *P*, the susceptible adult becomes infected and enters the latent phase of infection. If desired, the source of infection can then be determined by sampling from the relative contributions of each infectious contact to *P* (for instance, in this example, an infection is most likely to be transmitted from the household child, but all five infectious contacts have a finite probability of being identified as the source).

Disease natural history

We use the same influenza natural history model as used previously (4), which for completeness is recapitulated in Fig. 3. The main points are that the latent, incubation,

and contagious period durations are each sampled from discrete distributions, with mean periods of 1.2, 1.9, and 4.1 days, respectively. (The contagious period includes both the slight difference between latent and incubation periods, as well as the standard postincubation period when symptoms appear in 67% of infected people.) Any infectiousness that is not accompanied by overt symptoms (namely, the postlatent part of the incubation period, if any, and the 33% of infected people who never develop symptoms) is assumed to be half as great as the infectiousness of symptomatic individuals, reducing the transmission probability p_i by a factor of two.

As before (1-5), we also allow people who become ill to withdraw from all contact groups except their household, with an age-dependent withdrawal probability and distribution of the number of days of illness before withdrawal for influenza taken from Elveback *et al.* (1).

Pandemic influenza model parameterization

The potential pandemic influenza strain was assumed to have an age-dependent attack rate pattern between the historical 1957-1958 “Asian” influenza A (H2N2) (6) and 1968-1969 “Hong Kong” influenza A (H3N2) (7) pandemic strains (see Table 4). For fitting purposes, the attack rate pattern was calculated as an average of the final state of 500 independent communities that initially had 12 random infected people each (Fig. 4). As a baseline, the contact rates in households, small play groups, and large day care centers were taken from ref. 3, where an H2N2 strain was modeled. However, since this attack rate pattern hits school-age children particularly hard (see Table 4), these rates had to be reduced by about a factor of 3 (this is also evident in ref. 4, where a similar attack rate pattern was fit to a model specific for Thailand). The rest of the contact contribution was split between the remaining four contact groups (workgroups, household cluster, neighborhood, and community). Fine tuning to generate the contact probabilities shown in Table 3 was done by calculating the gradient vectors for the different age-dependent attack rates with respect to the contact rate parameters, which gives a linear approximation of the dependence of the attack rates as a function of the contact rates.

Although the fitting was merely done for isolated communities, we find that the national model has a very similar attack rate pattern (see Table 4).

2. Data Sources

The fundamental geographic unit in our model is the census tract, which is defined as a relatively stable geographic area with between 1,500 and 8,000 residents, with an optimum size of 4,000 people. In the 2000 Census, there were 65,443 census tracts containing 281,421,906 people in the U.S. (50 states and the District of Columbia), corresponding to an average population of 4,300 per tract (see Fig. 5 for the actual distribution of population sizes; www.census.gov/geo/www/tallies/tabgeo2k.html). We round off the population of each tract to the nearest 2,000 persons and populate each tract with the appropriate number of 2,000-person communities, each with households, schools, and other mixing groups as described above. In addition, several urban tracts have little or no residential population, but a large daytime worker population. We model these by communities comprised solely of work groups (in addition to the broad but weak community-level mixing), with an average of five 20-person work groups per each such community (corresponding to the average number of work groups in the suburban community model). In this way, we are able to realistically differentiate primarily residential tracts (with few, if any, work groups) from primarily urban ones (including some with few or even zero households). Each of the 180,492 model communities making up the national model is stochastically generated in an independent manner, so that no two communities within the nation are exactly identical.

Workplace tracts are chosen using the tract-to-tract worker flow data from the 2000 Census (CD-ROM special tabulation of Census 2000 data, available at www.census.gov/mp/www/spectab/stp64-webpage.html), which also provides the total number of working (and conversely, of unemployed) adults in each tract. The distribution of home-to-work commuter distances (measured from one tract center to another and zero if the home and work tracts are identical) is shown in Fig. 5. We note that these raw data refers to where individuals were working during the Census 2000 reference week

(generally the last week of March 2000), which is why a significant number of people (1.13 million, or 0.9% of the total workforce) were reported as working at locations 100 miles or more from their residence. We assume that such travel does not occur on a daily basis and instead place these individuals in a workgroup in their home tract. A related issue is the workers who were sick, on vacation, or otherwise absent from work during the reference week, estimated at 2% of used persons. Because both vacations and sick leave (withdrawal from workgroups) are included in the model, we compensate for these uncounted workers by multiplying each tract-to-tract worker flow total by 1.02.

The third source of data in our model captures the infrequent and irregular long-distance travel, such as business trips or vacations. We base this component on the 1995 American Travel Survey data available from the U.S. Department of Transportation, Bureau of Transportation Statistics*. The 1.00 billion person-trips (defined as 100 miles or longer each way, within the U.S.) among the 263 million residents at that date leads to an average of 3.8 trips per person, which we allocate according to the age group-specific data in Table 5. The average trip duration according to these data are 4.3 nights; we choose from a distribution between 0 and 11 nights according to the data in Table 5. For the present implementation, each trip destination is a random neighborhood within a random community (including workgroup-only communities), which results in a simple "gravity" model with no distance information. The destination community determines what types of contacts the traveler may have, in addition to the broad (but low-level) neighborhood and community-level mixing groups. During the daytime, the traveler may interact with his or her peers in play, school, or work groups if such contact groups exist at the destination tract; and at nighttime, the traveler may interact with a household and household cluster if traveling to one of the 78% of communities that are residential. In future work, we plan to incorporate a more sophisticated model of long-distance travel, which includes household and median destination income in determining trip frequencies, travel purposes (which affect both the choice of destinations and the relevant contact groups at the destination), and purpose- and distance-dependent distributions of trip durations (J. P. Newman, T.C.G., K.K., and C.A.M., unpublished work). As an additional step, the trips that are identified as air travel may be sampled from airline flight data,

capturing the long-distance travel component of disease transmission as realistically as possible.

3. Computational Implementation

To implement this computationally demanding model, we use the high-performance parallel molecular dynamics code "Scalable Parallel Short-range Molecular dynamics" (SPaSM) (8), written in C with message passing interface (MPI) communication. In recent years, this code has been used to model liquid- and solid-phase systems containing millions to billions (9) of atoms, yielding insights into such varied physical processes as dislocation dynamics (10), shock wave-induced plasticity (11), phase transformations (12) in metals, and fluid instabilities (13). The present epidemiological model is readily implemented in SPaSM (and presumably in similar particle-based codes) by replacing the C data structure for atoms (consisting of properties such as particle type, position, velocity, …) with one for persons (age, contact groups, immune system status, …), interatomic force field interactions with a social network and disease transmission model, and atomic classical mechanical trajectories with individual mobility rules (from residence to workplace on a regular basis and occasional long-range travel). A typical production run on 256 central processing units (CPUs) of a 2,048-processor Intel Xeon 2.4-GHz cluster with Myrinet interconnect took 8-12 h to complete a simulation of 180 days; depending on the disease parameters and amount of output, more CPU time was necessary in some instances. In all, ≈200 production runs were performed, amounting to ≈70 CPU years of computer time.

B. Scenario

We assume that the pandemic influenza strain is introduced into the U.S. via arriving international passengers. Furthermore, we assume that, by the time of this introduction, there is an ongoing worldwide pandemic, so there is no particular country or region that can be isolated (for instance, restricting arriving international flights from Southeast Asia). Consequently, we consider the 14 largest international airport gateways (U.S.

Dept. of Transportation, *U.S. International Air Passenger and Freight Statistics Report*, http://ostpxweb.dot.gov/aviation/usstatreport.htm) in the continental U.S. (see Table 6) and introduce a small number of infected individuals each day. We do so by choosing a random tract and community within each county listed in Table 6 and randomly infect between 0 and N individuals (chosen randomly from a uniform distribution) in that community. We take N proportional to the number of international arrivals at each airport, assuming 1-10 potential infecteds per 10,000 daily passengers. This represents a group of individuals, such as a family or business travelers, flying into the U.S. from the assumed pandemic that is raging worldwide throughout the 180-day simulation. For most simulations, we assume two infected persons per 10,000 daily international passengers, but the sensitivity to this choice and to other issues related to the seeding of infecteds is discussed below.

C. Basic Reproductive Number R_0

The value of R_0 was calculated for different transmission probabilities by three different methods, yielding the results summarized in Table 7. The first method was to average the number of secondary infections (omitting any tertiary infections) in 128,000 isolated communities that each had one random index case within the 2,000-person community population. (Using a smaller number of realizations, up to several thousand, led to statistical errors too great to determine R_0 within the desired ± 0.1 precision.) An example of this calculated distribution and the resulting average R_0 are shown for $P_{\text{trans}} = 0.12$ in Fig. 6, and these results are denoted “random index case” in Table 7.

The second method was similar, except that separate R_0 first were calculated for index cases belonging to each of the five age groups (see Table 7). The overall R_0 was then calculated as an average of these age group-dependent R_0 values, weighted by the age-dependent attack rate pattern for the respective transmission probability (referred to as the “attack rate pattern weighted index case” in Table 7). By doing this, the index case is more “typical” of those hit hardest by the outbreak and, as expected, this method slightly increases the value of R_0 (particularly for low R_0).

The last method is an approximation based on the slope of the cumulative number of cases (14) and allows also for a time dependence of the reproductive number $R(t) = 1 + \lambda\nu + f(1-f)(\lambda\nu)^2$, where ν is the sum of the latent and infectious periods, which for our model is 1.2 + 4.1 = 5.3 days, f is the relative duration of the latent period (i.e., 1.2 days/5.3 days for our model), and λ is the time derivative of the logarithm of the cumulative number of cases N, i.e., $\lambda = d[\ln(N)]/dt$. By seeding the 14 major international hubs with eight initial infected per 10,000 daily international passengers only at day 0 of the simulation, we calculated the basic reproductive number as a function of time (see Fig. 7). Although there are large oscillations at early times due to the larger statistical errors (from fewer cases), it is clearly noticeable that R is largest at early times and drops later. This is related to the fact that school children are particularly important spreaders in the initial stages of an influenza outbreak (see Fig. 7) due to their strong household and school interactions, which then naturally enhances R (see Table 7). After ≈30 days, the value for R stabilizes; in Table 7, we report (as "slope of cumulative number of cases") the average value between day 30 and the time when R begins to decline sharply, because large parts of the nation are already affected. This approximation is in good agreement with the value obtained by the second method, involving only single-community short-time simulations. We should note that the two former methods can give only an averaged static value as an estimator, whereas the latter method can give information about the time development of R. Here, early-time fluctuations and enhancement of the reproductive number clearly demonstrate the difficulty in measuring this quantity from available data in a real epidemic. Furthermore, this behavior is even more complicated by the spatiotemporal spread of the epidemic, which causes local variations of R in time. It should not be too surprising that the averaged static value of R_0 for the community and national models are similar, because, although the index case can interact with more than one community in the latter simulation, the number of effective interactions is the same, and in both cases, the populations are completely susceptible.

D. Intervention Strategies

1. Targeted antiviral prophylaxis (TAP)

Upon activation of a TAP program, symptomatic individuals and their close contacts are treated with antiviral drugs, until a possibly limited national stockpile has been exhausted. We assume that *X*% of symptomatic cases can be identified, and that 1 day after the onset of illness, the sick individual is treated therapeutically and prophylaxis offered to his/her close contacts. Of these, we assume that 100% of household, household cluster, preshool, and playgroup contact are identified and treated, and that *Y*% of workgroup and elementary, middle, and high school contacts are identified and treated. For the present work, we will focus on two cases: $X = Y = 60\%$ or 80%, and refer to these as "60% TAP" and "80% TAP," respectively.

As in other recent models of pandemic influenza (4), we use reported estimates of the antiviral efficacy for oseltamivir (15-19). Specifically, we assume that the antiviral efficacy for susceptibility $AVE_s = 0.30$, the antiviral efficacy for infectiousness $AVE_i = 0.62$, and the antiviral efficacy for illness given infection $AVE_d = 0.60$. For infected individuals, antiviral treatment reduces the infectious period by 1 day (whether or not the patient develops symptoms). Each course consists of 10 tablets, enough for 5 days of therapeutic treatment or 10 days of prophylaxis. If a person who is taking prophylactic course becomes ill, they complete their current course at the increased two tablets per day dosage. We assume that full antiviral efficacy is attained with the first tablet, and that there is no residual efficacy once the course has been completed.

2. Dynamic mass vaccination

Two major uncertainties in modeling any vaccination program are how effective the vaccine will be (because, even for endemic influenza, it is typically matched against a strain that is several months to 1 year old), and how quickly it can be produced, distributed, and result in an effective immune response. As yet, the efficacy of vaccination against a human-adapted avian strain is unknown. The immunogenicity of experimental vaccines has been measured; it has been found that a 4-fold increase in

antigen content above that of vaccines against human strains, and two vaccine doses, are required for a rise in antibody titer typically associated with protection (20). We assume that efficacy and immunogenicity are linearly related in our simulations. The second complication is related to timing. The time lag between vaccination and full effectiveness depends on the particular vaccine; for instance, a live attenuated vaccine may produce an antibody response within 1 day, whereas a killed vaccine may take 2 weeks. If multiple doses are required, the timescale can be significantly longer; for instance, two doses of a killed vaccine administered 4 weeks apart means that full efficacy may not be attained until 6 weeks after the initial dose. Rather than dealing with the specifics of any particular vaccine (including partial efficacy between the administration of the first and second doses), we simply combine this delay time with that for production and distribution and refer only to the date at which vaccination becomes effective, which may be either before or after the outbreak begins.[†]

We consider two alternative distribution strategies, either randomly throughout the entire eligible population or preferentially to children (with any remaining vaccine then distributed among adults). In either case, the eligible population consists of all individuals who have not been vaccinated and are not currently symptomatic. For simplicity, we consider only two alternative production scenarios, either assuming the early distribution of a low-efficacy (e.g., a poorly matched) vaccine or the delayed production of a higher-efficacy vaccine. The well matched vaccine is assumed to require two doses and to have a vaccine efficacy for susceptibility $VE_s = 0.70$ (with a reduced $VE_s = 0.50$ for the elderly, age 65+), and a vaccine efficacy for infectiousness $VE_i = 0.80$. The poorly matched vaccine has only $VE_s = 0.30$ (for all age groups) and $VE_i = 0.50$ and is assumed to require only a single dose [which would not be the case for an avian influenza A (H5N1) virus-based vaccine, for instance (20)]. It is assumed that early production of the poorly matched vaccine allows for a vaccination program before the outbreak, resulting in a prior coverage of some fraction of the population (again, either uniformly or preferentially to children). For either vaccine, we assume a constant production and distribution rate of 4, 10, or 20 million doses per week nationwide, starting as soon as 2

months before the first introduction, to as late as 2 months after the first introduction. The total production is also assumed to be limited to 50, 100, 250, or 400 million doses.

3. School closure

Upon recognition of a pandemic strain in the U.S., one of the likely mitigation strategies is the closure of schools [U. S. Department of Health and Human Services (HHS) Pandemic Influenza Response and Preparedness Plan, www.hhs.gov/pandemicflu/plan]. We assume this involves a total nationwide closure all of the school-related mixing groups in our model, and that this closure remains in effect for the duration of the pandemic. The affected mixing groups are the regular preschool-age playgroups, preschools, and elementary, middle, and high schools. All other contact rates remain unchanged.

4. Social distancing

As a result of either a formal quarantine program or voluntary changes in social and hygienic behavior in the event of a widespread pandemic, it is likely that normal contact behavior will be affected in times of crisis. Although this alteration is difficult to predict in advance, it almost surely will involve an increased tendency to remain at home rather than in large public places. To approximate this behavioral modification, we assume that the contact rates are cut in half for the community, neighborhood, work group, school, preschool, and playgroup mixing groups; household contact rates are doubled; and household cluster contact rates remain unchanged. As with the other mitigation strategies, it is assumed that this alteration in normal behavior occurs nationwide and lasts throughout the remainder of the epidemic.

5. Reductions in travel

We consider reductions in both of our travel components: the daily workplace travel and irregular long-range travel. The first may be curtailed by voluntary increases in

telecommuting or, in extreme cases, by a nationwide work stoppage (with the exception of health care and emergency personnel, as described below). Reductions in long-range travel may also range from a component of the natural social distancing tendency to an imposed quarantine or travel restriction program. We assume that long-range travel may be reduced to as little as 1% of the normal number of trips.

E. Sensitivity Analyses

In this section, we explore sensitivity to various components of the model, including how the epidemic is introduced, delays in implementing intervention strategies, and the assumed effectiveness (or public compliance) with each intervention. Although each of these variations naturally leads to quantitative changes in the precise quantitative results, the basic conclusions presented in the main text are all relatively insensitive to any of the variations that we have explored.

1. Stochastic variability

Because the model is inherently stochastic (including mock population generation, introduction of new infecteds, daily disease transmission, and intervention strategy components), in theory, we need to run several realizations for each scenario, each with a different initial random seed. (For parallel runs, each processor uses a different seed to avoid having exactly identical communities anywhere in the simulation.) However, we have observed that between the large degree of spatial averaging over the 180,000 communities making up the national model and the daily introduction of new infecteds, which demands a robust intervention strategy, there is almost no variation in either the nationwide epidemic curve or final attack rate but only subtle differences in the specific timing and geographic details of the epidemic spread. For instance, we compared eight different baseline (no intervention) realizations for $R_0 = 1.6$, differing only by the initial random number generator seeds. All eight simulations give nearly identical final attack rates (four simulations give 32.62%, three give 32.63%, and one gives 32.64%), with the

epidemic peaking between 115 and 120 days after the first U.S. introduction in all realizations.

2. Seeding of the epidemic

To investigate sensitivity against different introductions of infected individuals (size and spatial distribution of these individuals, as well as dynamic vs. static seeding), different amounts of travel, as well as different random seeds (i.e., stochastic behavior), a series of runs with no intervention is presented for $R_0 = 1.9$ (similar results were obtained for other choices of R_0). Fig. 8 shows the sensitivity to the dynamic seeding rate, from one to four infected persons per 10,000 international passengers each day at the major 14 U.S. international air hubs, as well as a static seeding of eight infecteds per 10,000 international passengers only at the beginning of the simulation (corresponding to 76 infected people introduced on day 0). It can be clearly seen that the size of the seed and the effect of static vs. dynamic seeding only shifts the epidemic curves without affecting the shape or overall attack rate (43.53% in all cases).

The static seeding on day 0 is further explored in Fig. 8, comparing the dispersed introduction of 76 infected people at 14 airports with 40 infected people all localized either in Los Angeles County or New York County. Once again, the different seeding changes only the timescale of the pandemic outbreak but not the shape or magnitude of the epidemic curve. Although nationally averaged measures such as this (or others, including the number of antiviral courses required for a TAP intervention) do not depend on the size or distribution of the seeding, the precise details of the spatiotemporal evolution do. This is illustrated for the Los Angeles and New York County seeds in Fig. 9, showing that very different geographic spreads can yield virtually identical national epidemic curves (Fig. 8).

Fig. 8 also shows the effect of drastic reductions in long-range travel, to only 1-10% of the normal levels during the entire 180 days of the simulation, for an initial introduction in Los Angeles County. Here it can be seen that the width of the epidemic curve widens

and the peak shifts to later times, both useful factors when considering the demand upon the health care system and resource allocation. Even though these reductions reflect nearly a complete halt of nonessential travel (other than local commuter travel to workplaces), with only 1–10% of leakage or essential emergency travel, the total attack rate after 180 days is unchanged. Also shown in Fig. 8 is the epidemic curve for a simulation for which a travel reduction to 1% of normal levels is imposed only after the pandemic alert threshold, which is reached on day 38 in this case. Here one can see that the (already marginal) effectiveness of travel restrictions is reduced even further if the virus is given any time to spread, because it may introduce many small pockets of infection that are able to develop despite the draconian measures.

3. Vaccine production, distribution, and effectiveness delays

Because an intervention strategy of vaccination alone is unlikely to ever succeed for $R_0 >$ 1.9 (see Table 2), we show in Table 8 the effectiveness of different production rates, limits, and starting dates for $R_0 = 1.6$. In addition to the necessity for both high production rates and limits (which are more important than the exact starting dates), we find a clear advantage to a preferential vaccination of children, as has been suggested recently (21). Perhaps more surprisingly, we find that a more widespread vaccination coverage with a lower efficacy is decidedly more effective than a higher-efficacy vaccination of half as many people (also shown in Fig. 10), even before taking into account the additional 4-6 weeks, which may be required to elicit a strong immune response from a two-course vaccine program.

4. Targeted antiviral prophylaxis: Delays in policy implementation and patient diagnosis

There are two timescales that may affect the effectiveness of a TAP program. The first of these is at the population-wide public health scale, involving the time required to recognize that a nascent pandemic outbreak is underway and implement a public health response. The second timescale is at the individual level, involving the time it takes from

the first appearance of symptoms, before the person visits their doctor or urgent care center, to a correct diagnosis and prescription of antiviral treatment for the patient, to prescription and delivery of prophylactic courses to that patient's close contacts. In Fig. 10, we show (for $R_0 = 1.9$) the increasing number of antiviral courses required and the increasing attack rate, the longer it takes to initiate a TAP program. (Day 0 on this axis refers to a policy that is in place even before the first introduction of a pandemic influenza strain into the U.S.) Clearly, an early intervention has benefits by reducing both the incidence rate and the demand on the limited antiviral supply.

In the event (or even the threat) of a nascent pandemic, a major challenge will be to correctly distinguish patients with the pandemic influenza strain from the larger number of individuals with general respiratory symptoms (possibly, but not limited to, those caused by a conventional strain of influenza). In a typical flu season, only 10-30% of patients presenting flu-like symptoms actually test positively for influenza; were antiviral treatment and prophylaxis of close contacts to be offered based on symptoms alone, a limited stockpile could be rapidly exhausted. The TAP interventions modeled here assume the availability of a large number of rapid test kits and testing locations (e.g., workplaces and convention centers, in addition to health care offices). With such a capability, positive identification and antiviral distribution could be achieved within 1 day of the first appearance of symptoms. To examine the sensitivity to this assumption, we have also carried out a few simulations with a 2-day lag between the appearance of symptoms and the antiviral treatment and close-contact prophylaxis, as shown in Table 8 for an 80% TAP intervention.

Finally, Table 9 compares the 60% and 80% TAP programs, indicating the need for both early intervention (as seen in Fig. 10) and for a high identification rate. For instance, a nascent pandemic outbreak with $R_0 = 1.9$ can be contained with only 27 million courses if an 80% TAP strategy is implemented as late as 7 days after the pandemic alert threshold (starting 3 days earlier, at 4 days after the alert, requires only 20 million courses), which is within the planned national antiviral stockpile. However, a 60% TAP program initiated at the same time is moderately successful but consumes 182 million courses, currently a

prohibitively large supply of antivirals. We also find that it is important to be able to ascertain close contacts beyond the household, including school and work group peers and household clusters; 60% TAP with only prophylaxis of the household uses few courses but with a significantly higher total attack rate.

5. Social distancing strategies, including travel restrictions

As discussed in the main text, all of the social distancing strategies we have examined serve only to slow down the epidemic spread but are ineffective at reducing the overall attack rate (shown in Table 9 for a combination of all social distancing measures). The example of long-distance travel restrictions is discussed in *Sensitivity Analyses* 2 for the static one-time introduction of infecteds into Los Angeles (Fig. 8) and is illustrated in Fig. 11 and Table 9 for the usual daily introduction through airport hubs, for $R_0 = 1.9$ and 2.4. Although the final national attack rates are virtually unaffected, drastic reductions in long-range travel to 1–10% of the normal rates clearly spreads out the pandemic into two waves: after an initial peak in the sites of introduction (which occurs at the same time as the baseline case without any intervention), a secondary peak as long as 2 months later corresponds to the nationwide spread. In addition to buying time for vaccine production and other mitigation strategies to be used, the distinct spatial evolution may aid the health care response, because resources can gradually be shifted from the sites of initial introduction to the secondary sites in the rest of the country.

6. Combined mitigation strategies

Table 10 shows results for several combinations of the mitigation strategies, both therapeutic (vaccines and/or antivirals) and social distancing (including school closure and travel restrictions). The key ingredients of any response plan seem to include both TAP and school closure if a pandemic with potential R_0 as high as 2.4 is to be avoided, although one of these measures (but not both) may be omitted if $R_0 < 2.1$.

F. Simulation Movies

All movies are single realizations for $R_0 = 1.9$, with the standard daily introduction of infected people through 14 major international airports, as described above. As in Fig. 1 of the main text, each census tract is shown as a dot colored according to the current prevalence, on a logarithmic scale from green for 0.03% or fewer ill people per capita, to red for 3% or greater. The corresponding epidemic curves (averaged over the entire nation) are also shown.

*Available at www.bts.gov/publications/national_transportation_statistics. More recent (2001) data define long-distance as trips 50 miles or greater of which nearly half (47.7%) are <200 miles roundtrip and only 7.4% are via airplane (of 9.2 such trips per person each year). Although the 1995 definition of long-distance travel as 100 miles or greater is still dominated by automobile travel (81.3%), the mean roundtrip distance of 826 miles is much more suitable for the distance-independent gravity model used here.

†This leads to a minor inconsistency, in that the eligible population is determined at the date at which full effectiveness is reached and not at the earlier date of vaccination. However, this study is concerned with strategies to minimize the number of infected individuals, in which case the number of new symptomatic cases (who become ineligible for vaccination) between these two dates is negligible.

1. Elveback, L. R., Fox, J. P., Ackerman, E., Langworthy, A., Boyd, M. & Gatewood, L. (1976) *Am. J. Epidemiol.* **103,** 152-165.

2. Halloran, M. E., Longini, I. M., Cowart, D. M. & Nizam, A. (2002) *Vaccine* **20,** 3254-3262.

3. Longini, I. M., Halloran, M. E., Nizam, A. & Yang, Y. (2004) *Am. J. Epidemiol.* **159,** 623-633.

4. Longini, I. M., Nizam, A., Xu, S., Ungchusak, K., Hanshaoworakul, W., Cummings, D. A. T. & Halloran, M. E. (2005) *Science* **309,** 1083-1087.

5. Halloran, M. E., Longini, I. M., Nizam, A. & Yang, Y. (2002) *Science* **298,** 1428-1432.

6. Kilbourne, E. D. (1975) *The Influenza Viruses and Influenza* (Academic, New York).

7. Rvachev, L. A. & Longini, I. M. (1985) *Math. Biosci.* **75,** 3.

8. Beazley, D. M.& Lomdahl, P. S. (1994) *Parallel Comput.* **20,** 173.

9. Kadau, K., Germann, T. C. & Lomdahl, P. S. (2004) *Int. J. Modern Phys. C* **15,** 193.

10. Zhou, S. J., Preston, D. L., Lomdahl, P. S.& Beazley, D. M. (1998) *Science* **279,** 1525-1527.

11. Holian, B. L. & Lomdahl, P. S. (1998) *Science* **280,** 2085-2088.

12. Kadau, K., Germann, T. C., Lomdahl, P. S. & Holian, B. L. (2002) *Science* **296,** 1681-1684.

13. Kadau, K., Germann, T. C., Hadjiconstantinou, N. G., Lomdahl, P. S., Dimonte, G., Holian, B. L. & Alder, B. J. (2004) *Proc. Natl. Acad. Sci. USA* **101,** 5851-5855.

14. Lipsitch, M., Cohen, T., Cooper, B., Robins, J. M., Ma, S., James, L., Gopalakrishna, G., Chew, S. K., Tan, C. C., Samore, M. H., *et al.* (2003) *Science* **300,** 1966-1970.

15. Hayden, F. G. & Aoki, F. Y. (1999) in *Antimicrobial Therapy and Vaccines*, eds. Yu, V., Meigan, T. & Barriere, S. (Williams & Wilkins, Baltimore), pp. 1344-1365.

16. Hayden, F. G., Gubareva, L. V., Monto, A. S., Klein, T. C., Elliott, M. J., Hammond, J. M., Sharp, S. J. & Ossi, M. J. (2000) *N. Engl. J. Med.* **343,** 1282-1289.

17. Monto, A. S., Robinson, D. P., Herlocher, M. L., Hinson, J. M., Elliott, M. J. & Crisp, A. (1999) *J. Am. Med. Assoc.* **282,** 31.

18. Welliver, R., Monto, A. S., Carewicz, O., Schatteman, E., Hassman, M., Hedrick, J., Jackson, H. C., Huson, L., Ward, P. & Oxford, J. S. (2001) *J. Am. Med. Assoc.* **285,** 748.

19. Yang, Y., Longini, I. M. & Halloran, M. E. (2006) *Appl. Stat.*, in press.

20. Altman, L. C. (August 7, 2005) *N.Y. Times*, p. 1.

21. Longini, I. M. & Halloran, M. E. (2005) *Am. J. Epidemiol.* **161,** 303.