

Examining Ontario's universal influenza immunization program with a multi-strain dynamic model



E.W. Thommes^{a,b,*}, A. Chit^{c,d}, G.C. Meier^e, C.T. Bauch^{f,b}

^a Medical Division, GlaxoSmithKline Inc., 7333 Mississauga Road, Mississauga, ON, Canada L5N 6L4

^b Department of Mathematics & Statistics, University of Guelph, 50 Stone Road East, Guelph, ON, Canada N1G 2W1

^c Sanofi Pasteur, 1755 Steeles Avenue West, Toronto, ON, Canada M2R 3T4

^d Leslie Dan Faculty of Pharmacy, University of Toronto, 144 College Street, Toronto, ON, Canada M5S 3M2

^e GlaxoSmithKline Inc., Avenue Fleming 20, 1300 Wavre, Belgium

^f Department of Applied Mathematics, University of Waterloo, 200 University Avenue West, Waterloo, ON, Canada N2L 3G1

ARTICLE INFO

Article history:

Received 12 February 2014

Received in revised form 13 April 2014

Accepted 2 June 2014

Available online 10 July 2014

Keywords:

Influenza

Dynamic transmission model

Universal vaccination

TIV

Ontario

Canada

ABSTRACT

Seasonal influenza imposes a significant worldwide health burden each year. Mathematical models help us to understand how changes in vaccination affect this burden. Here, we develop a new dynamic transmission model which directly tracks the four dominant seasonal influenza strains/lineages, and use it to retrospectively examine the impact of the switch from a targeted to a universal influenza immunization program (UIIP) in the Canadian province of Ontario in 2000. According to our model results, averaged over the first four seasons post-UIIP, the rates of influenza-associated health outcomes in Ontario were reduced to about half of their pre-UIIP values. This is conservative compared to the results of a study estimating the UIIP impact from administrative data, though that study finds age-specific trends similar to those presented here. The strain interaction in our model, together with its flexible parameter calibration scheme, make it readily extensible to studying scenarios beyond the one explored here.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

Seasonal influenza is responsible for a significant health burden each year, infecting roughly a tenth of the world's population (e.g. [1]). Large-scale changes in a population's vaccination patterns provide a good "natural laboratory" for better understanding the dynamics of an infectious disease such as influenza, and for testing epidemiological models. Just such a scenario took place in the province of Ontario, Canada, which in 2000 initiated the world's first large-scale universal influenza immunization program (UIIP), whereby influenza vaccination was provided for free to all residents. Subsequently, Kwong et al. [2] utilized provincial administrative data to study the impact of the UIIP on influenza-related health outcomes. Here, our objective was to develop a general-purpose seasonal influenza model, use it to simulate Ontario's UIIP adoption, and to test the results against those of [2].

Several key factors determined the design of our model. First and foremost, with a basic reproduction number R_0 estimated to range

from 1.6 to 3.9 [3,4], the infectiousness of influenza is sufficiently low that herd immunity is important even at modest levels of vaccine coverage within a population; self-consistently accounting for herd immunity requires a dynamic transmission model. Second, we strove to make the model robust and flexible enough to be applied to a wide range of scenarios beyond the one explored here. Public health systems worldwide have a wide array of existing (e.g. inactivated, live-attenuated, adjuvanted, unadjuvanted) and new (e.g. quadrivalent, cell-cultured) seasonal vaccines to choose from, and with numerous additional influenza vaccines on the horizon [5], the selection is likely to become wider still. Mathematical models can serve as powerful tools to assist policymakers in the optimal adoption of these technologies.

We wanted to implement a transmission model sophisticated enough to reproduce the key dynamics of seasonal influenza – herd immunity, strain interaction, waning immunity, dependence on population contact patterns – while at the same time simple enough to be straightforwardly calibrated to real-world data. Accordingly, we chose a Susceptible–Infected–Recovered–Vaccinated compartmental model as our basic paradigm. We extended this approach to explicitly model the four dominant strains/lineages of seasonal influenza, with cross-protection where relevant. To our knowledge,

* Corresponding author at: Medical Division, GlaxoSmithKline Inc., 7333 Mississauga Road, Mississauga, ON, Canada L5N 6L4. Tel.: +1 905 819 3480.

E-mail address: edward.w.thommes@gsk.com (E.W. Thommes).

previous age-structured influenza models have included a maximum of three strains, and no three-strain models have considered cross-protection [6,7]. Furthermore, we equipped the model with an approximate Bayesian computation (ABC) [8] parameter fitting scheme which is robust, flexible, and directly propagates the effect of parameter uncertainty to the model results.

Explicitly modeling both lineages of influenza B introduces an added level of complexity in how vaccination is implemented within the model. The reason is that current seasonal trivalent influenza vaccines (TIVs) include both dominant circulating subtypes of influenza A (H1N1, H3N2), but only one of the two circulating influenza B lineages (Victoria, Yamagata). Hence, there is the possibility of influenza B vaccine lineage mismatches: For example, B/Victoria is chosen for inclusion in a given season's TIV, but B/Yamagata ends being the dominant B lineage circulating that season. Historically, such a mismatch has occurred in roughly half of all seasons. Thus, when modeling the efficacy of TIV against influenza B over multiple seasons, care must be taken that the match/mismatch proportion reflects what happens in real life. This becomes especially important if one wishes to use the model for a comparison of TIV with quadrivalent (QIV) vaccines, which includes both B lineages; having too low or too high a rate of B lineage match for the TIV will under- or over-estimate, respectively, the effectiveness of TIV relative to QIV.

2. Materials and methods

2.1. Age structure and demographics

Our model was age-stratified [6,9], enabling us to include age-dependent contact patterns in calculating the force of infection. Births, deaths and aging were implemented; age-specific birth and death rates can be specified for the population, and these can be made to vary on a year-by-year basis. More detail on age structure and demographics in the model is given in A.1.

2.2. Epidemiology

The model explicitly tracks the two dominant A strains (H1N1, H3N2), and the two B lineages (Victoria, Yamagata) comprising seasonal influenza. We assumed that A and B cross-protection, both natural and vaccine-conferred, is negligible, and only allowed for the possibility of pairwise cross-protection between the two A strains, and between the two B lineages, respectively. We modeled seasonality in transmission as a sinusoidal variation in the force of infection, with a period of one year.

Immunity to influenza strains conferred by infection is only temporary, due to a combination of declining antibody levels and the continual antigenic drift of the influenza virus; see e.g. [10]. We assumed that individuals lose their immunity to a given strain k at a constant rate θ_k [6,11,12]. The inverse, $1/\theta_k$, is then the mean duration of natural immunity. Details of the transmission model are given in Appendix A.2.

2.3. Vaccination

Vaccination occurs over a finite time window, the yearly start time and length of which are input parameters. The vaccine is characterized by its efficacies against each of the circulating strains/lineages of influenza. TIV contains two A strains and only one B lineage, thus each season one of the B lineages is matched by the vaccine, while the other is mismatched. Hence, for B, the model requires both a matched and a mismatched efficacy as input. How well a given season's vaccine protects against influenza B is determined by the proportion of matched and mismatched B lineages that are in circulation that season; in a dynamic model, as in real

life, this proportion is not known *a priori*. In light of this we made the accuracy of the choice a fitted parameter which can vary anywhere from a random choice to a "best guess". Details are given in Appendix A.3.

Vaccine-conferred immunity also wanes due to declining antibody levels and antigenic drift, and this process was implemented in the model analogously to the waning of natural, infection-conferred immunity, i.e. via a constant rate of immunity loss Θ_k for strain k . Since vaccine-conferred immunity is generally thought to be less robust than natural immunity, it was considered important that the model allow the two to wane at different rates. Details are given in Appendix A.3.

3. Model calibration

Since our understanding of the natural history of influenza is still far from complete, choosing realistic values for some of the parameters in the model poses a significant challenge. We adopted an approximate Bayesian computation (ABC) scheme [8] analogous to one used previously for fitting human papillomavirus models [13,14]:

- (1) Each model input parameter to be fitted is given a uniform prior distribution between some minimum and maximum;
- (2) Model outputs (model summary statistics) to be fitted are chosen. For each, a minimum and maximum value are chosen, thus defining an allowable target interval.
- (3) Sets of parameters are drawn from the prior distributions. Latin hypercube sampling [15] is used in order to obtain a more even coverage of the parameter space with fewer samples than simple random sampling would yield.
- (4) Each set of parameters drawn is used in one model run. The posterior parameter distribution consists of all sets of parameters which result in a model run in which all outputs simultaneously satisfy their respective allowable target ranges.

In this way, even though we may know little or nothing about the true value of an *individual* parameter (e.g. natural cross-protection; see Table 1), the posterior parameter distribution will consist only of *combinations of parameters that yield reasonable model outputs*. We restricted our fitting to natural history parameters, that is, parameters intrinsic to influenza itself, rather than parameters describing the population (e.g. birth rate, vaccine uptake, contact patterns), since natural history parameters are in general much more uncertain. Also, natural history parameters can be considered largely independent of the population, thus a posterior distribution obtained for one setting can be applied to another, as long as both populations possess a broadly similar lifestyle and average health status (e.g. if both are developed countries). This allowed us to exploit the availability of more suitable calibration target data for the US rather than for Canada; we conducted our natural history calibration in the former setting and applied the results to the latter. A total of 5000 fitting simulations were run, yielding 181 posterior (i.e. accepted) parameter sets. Details are given in Appendix B.

4. Modeling Ontario's transition to universal influenza immunization

Using the posterior parameter sets yielded by our model calibration above, we performed simulations of Ontario's adoption of a universal influenza immunization program (UIIP). For each simulation, one of the 181 posterior parameter sets was drawn at random, together with a new random seed (thus preventing two draws of the same parameter set from producing identical simulation results). Sets of 100, 500 and 1000 such

Table 1
Baseline epidemiology- and vaccine-related model input parameters.

Input parameter	Definition	Value or prior range	Source
$\tau_{A,nat}$	Mean duration of natural immunity to influenza A	1–6 yrs	Trial and error; posterior density found to be zero outside this range
$\tau_{B,nat}$	Mean duration of natural immunity to influenza B	4–35 yrs	Trial and error; posterior density found to be zero outside this range
$\Xi_{A,nat}$	Natural cross-protection between A strains (H1N1, H3N2)	0–1	No data on H1N1–H3N2 natural cross-protection was found, thus maximum possible range used
$\Xi_{B,nat}$	Natural cross-protection between B lineages	0.44–1	TIV cross-protection against mismatched B lineage (see below) assumed as lower bound for natural cross-protection
β'	Probability of infection per unit of contact (=1 min) with an infected person	2.8695e – 4	Chosen to yield a mean $R_0=1.3$, which together with $A_{seasonal}$ (below) yields a seasonal peak $R_0 \approx 1.9$ [19].
$A_{seasonal}$	amplitude of seasonal variation factor $F_{seasonal}$ multiplying the force of infection	0.43	[6]
P_{bBg}	probability of using the “best B-lineage guess” (Appendix A.3) to select the lineage to include in TIV. 1 = always use best guess, 0.5 = random choice, 0 = always use opposite of best guess.	0–1	Maximum possible range
t_{start}	Day of the year on which new flu season starts, and previous season ends	Day 213 (August 1)	Chosen as middle of summer, when Northern Hemisphere seasonal influenza activity is at low point
$t_{max,seasonal}$	Day of the year on which $F_{seasonal}$ and thus the force of infection, takes on its maximum value	Day 365 (December 31)	Chosen to produce seasonal influenza epidemic peaks early each year, consistent with Northern Hemisphere seasonal influenza surveillance data; see e.g. Fig. 1
t_{vacc}	Day of the year on which vaccination campaign begins	day 274 (October 1)	In the US, influenza vaccination typically commences in October and finishes at the end of the year (http://www.cdc.gov/flu/professionals/vaccination/vax-summary.htm)
Δt_{vacc}	Duration (in days) of vaccination campaign	92 days	see t_{vacc} above
$F_{symptomatic}$	The fraction of influenza cases which are symptomatic	0.669	[32]
f_{asympt}	Infectiousness factor for asymptomatic cases (1 = no reduction vs symptomatic; 0 = asymptomatic cases cannot infect others)	1.0	assumption (any intrinsic reduction in infectiousness of asymptomatic relative to symptomatic cases will be at least partially countered by their greater likelihood to have contact with others)
Δt_{flu}	Mean latent + infectious period	4 days	[6] and references therein
ϵ_A	TIV efficacy against Influenza A	0–17 yrs: 0.59 18–64 yrs: 0.61 65+ yrs: 0.58	Taken from a meta-analysis of vaccine efficacy against matched and mismatched strains/lineages in adults [33]. Reduction of efficacy in children and the elderly is extrapolated using previous meta-analyses of vaccine efficacy [34,35] (which averaged over matches and mismatches)
$\epsilon_{B,match}$	TIV efficacy against the included B lineage	<5 yrs: 0.66 5–49 yrs: 0.77 50–64 yrs: 0.73 65–74 yrs: 0.69 75+ yrs: 0.66	
$\epsilon_{B,mismatch}$	TIV efficacy against the non-included B lineage	<5 yrs: 0.44 5–49 yrs: 0.52 50–64 yrs: 0.49 65–74 yrs: 0.47 75+ yrs: 0.44	
τ_{nat}	Mean duration of vaccine-conferred immunity	1 yr (all ages, strains, lineages)	Conservative assumption, based on a study [36] showing vaccine efficacy in the elderly dropping to very low levels after one season

simulations were performed; 1000 were deemed sufficient, since results (mean and 95% CI) varied negligibly between the 500- and 1000-simulation set. Ontario demographic data (population distribution, birth and all-cause mortality rates) were taken from Statistics Canada's CANSIM online socioeconomic database, <http://www5.statcan.gc.ca/cansim/>. Since our objective was to compare these simulations to the results of Kwong et al. [2], our simulations covered the same period as that study, 1997–2004 (three seasons pre-UIIP, 4 seasons post-UIIP). Vaccine uptake rates used are given in Table 2. Vaccine uptake rates for ages 12 years

and up during these seasons were taken from Kwong et al. [2]; vaccine uptake in lower ages was extrapolated. Details are given in Appendix D.1.

The study of Kwong et al. [2] reports rates of influenza-associated events (mortality, hospitalization, emergency department use and visits to doctors' offices), whereas our model calculates the rate of influenza cases themselves, in other words the influenza infection attack rate (AR). Direct comparison to these event rates thus requires knowing the (age-dependent) probabilities with which an influenza case leads to each of the events,

Table 2
Age- and time-stratified Ontario vaccine uptake, taken from Refs. [2,39].

Year	Age (yrs) 0 (%)	1 (%)	2–11 (%)	12–19 (%)	20–49 (%)	50–64 (%)	65–74 (%)	85+ (%)
1985	1.03	2.05	16.00	16.00	8.00	21.00	54.00	70.00
1999	1.03	2.05	16.00	16.00	8.00	21.00	54.00	70.00
2000	1.03	2.05	29.00	29.00	27.00	42.00	69.00	79.00
2001	1.03	2.05	29.00	29.00	27.00	42.00	69.00	79.00
2003	5.93	11.85	28.00	28.00	23.00	45.00	71.00	80.00
2005	6.55	13.10	37.00	37.00	30.00	54.00	73.00	84.00




Fig. 1. A random selection of accepted fitting simulations, showing weekly incidence of influenza A (blue) and B (red). For comparison, WHO-NREVSS US surveillance data is shown in the bottom panel. Though simulation and surveillance data look qualitatively similar, note the thousand-fold difference in scale; whereas the model simulates all cases of influenza in the US, real-world surveillance captures only a small fraction. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

as calculated, e.g. in Molinari et al. [16] for the US. For example, using their notation, $P(\text{Hosp}/\text{flu})$ is the probability that a person will be hospitalized as a result of contracting symptomatic influenza. We were unable to identify a Canadian-specific source for such probabilities, and therefore used the US results of Ref. [16]. As explained in Appendix D.2, any differences between US and Canadian outcomes probabilities ought to have only modest impact on our results, since we are looking at the *relative* rates of outcomes going from pre- to post-UIIP Ontario, and since US and Canadian demographics and health status are overall similar.

4.1. Results

Fig. 2 shows a comparison of our model results for the relative rates of outcomes, $RR = R_{\text{post-UIIP}}/R_{\text{pre-UIIP}}$, (where % reduction = $[1 - RR] \times 100\%$) to those of [2]. Across all ages, the model RR point estimates for all four outcomes are just above 0.5 (i.e. a 50% reduction in outcomes), whereas those of [2] range from about 0.2 to 0.3 (i.e. a 70–80% reduction in outcomes). The 95% confidence intervals do not overlap for all but mortality (and there only barely); clearly, there is a statistically significant difference between the model results and those of Ref. [2]. However, an additional analysis step was performed by [2], in order to control for any systematic changes across Canada during the study period: They calculated *normalized RRs* by taking the ratio of the post- versus pre-2000 outcomes rates RRs for Ontario, to the post- versus pre-2000 RRs in the rest of Canada (where no universal immunization programs existed throughout the study period). We performed an analogous normalization by rerunning each of our simulations using the averaged vaccine uptake rates for the rest of Canada, then for each pair of simulations taking the ratio of the Ontario RR to the “rest-of-Canada” RR. Details are given in Appendix D.1. The results are shown in Fig. 3. Here, the overall agreement is better: For the RRs taken across all ages, the 95% CIs of the model results include Kwong et al.’s [2] point estimates (they did not calculate CIs for this quantity) for all outcomes. Agreement is generally better at high ages, whereas below age 65, the results of [2] all fall near the lower edge of our model RRs.

5. Discussion

Overall, our base case range of model results for the reduction in influenza-associated outcomes resulting from the introduction of Ontario’s universal influenza immunization program (UIIP) were consistent with the findings of Kwong et al. [2]. At the same time, model results were on average conservative, in the sense of showing smaller point estimates of post-UIIP reductions than [2] across all outcomes. Normalized to the rest of Canada, our simulations showed, across all ages, mean post- versus pre-UIIP relative rates (RRs) of outcomes ranging from about 0.7 (GP visits) to 0.8 (deaths) (see Fig. 3). In comparison, [2] found normalized RRs ranging from about 0.4 (GP visits) to 0.6 (deaths).

We designed the scheme for calibrating influenza natural history parameters of our model to require relatively few assumptions. One assumption we did make is that, for all ages, vaccine-conferred immunity only lasts one year on average. Sensitivity analysis showed (see D.3) that a longer duration of vaccinated immunity, as used in previous modeling studies [6,7], can lead to greater post-UIIP reductions, more similar to those found by Ref. [2]. This may suggest that a one-year duration for vaccine-conferred immunity across all ages is an overly conservative value.

As part of the calibration, we also found it necessary to impose a constraint ensuring that the basic reproduction number R_0 does not take on unrealistic values. However, this leads to the question of what constitutes a realistic R_0 for seasonal influenza. Since (unlike for pandemics) only part of the population starts out susceptible, one can directly calculate only the *effective* reproduction number, usually denoted R or R_p [17,18]. Our adopted value of R_0 (seasonal maximum 1.9) is taken from a maximum-likelihood fit to influenza surveillance data over multiple seasons and settings [19]. However, this study itself used a dynamic model as part of the fitting process, hence we are in effect taking output from another model as input to our model. In future work, we will explore how to refine the calibration process to render this constraint unnecessary.

Model uncertainty arises not only from the natural history parameters. Another limitation of our simulations is the extrapolation needed to obtain vaccine uptake for ages 2 to 11 in Ontario,




Fig. 2. Un-normalized model results (red) for the age-stratified post- versus pre-2000 (=beginning of universal influenza immunization program, UIIP) fractional changes (relative rates, RR) in influenza-associated outcomes. Filled circles show point estimates, and error bars show 95% CIs. The results obtained by Kwong et al. [2] are shown for comparison (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. Normalized relative rates: Model results (red) for the post- versus pre-2000 relative rates (RRs) of influenza-associated outcomes in Ontario, divided by the post- vs. pre-2000 RRs in the rest of Canada. Filled circles show point estimates, and error bars show 95% CIs. The results obtained by Kwong et al. [2] are shown for comparison (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and ages 0 to 11 in the rest of Canada. As demonstrated in the sensitivity analysis (Appendix D.3), inaccuracies in the vaccine uptake in this age group alone have the potential to substantially affect model results. Incidentally, this echoes previous findings (e.g. [20–22]) that vaccinating children can have a disproportionately large impact on program effectiveness. The fact that the normalized RRs produced by the model most strongly exceed those of [2] at younger ages may indicate that our extrapolation underestimated the post-2000 uptake in Ontario, or overestimated it in the rest of Canada, or a combination of the two.

Another limitation of our model can be seen when comparing simulated and real weekly incidence in Fig. 1: In real life, the timing of seasonal epidemic peaks is more variable than it is in the model results. The regularity of the latter is due to our method (commonly used also in previous models, e.g. [6,7]) of producing seasonality, namely multiplying the force of infection by a sinusoidal factor of constant period and phase. Since the effectiveness of a vaccination program will be reduced if it is still underway when that season's epidemic commences, any systematic difference between the timing of real and simulated epidemics is potentially problematic. However, in Fig. 1, variations notwithstanding, all surveillance data epidemic peaks except for the 2014–2015 season occur principally in the new year, by which time vaccination programs have generally wrapped up; it thus seems unlikely that these variations strongly affected vaccine effectiveness. Nevertheless, adding a realistic representation of the variability in the timing of influenza peaks, and also in the timing and duration of vaccination programs, would constitute a worthwhile future improvement to the model.

A structural limitation of our model is that it stratifies the population only by age. Hence it captures age-specific contact patterns within the population, but does not capture any population heterogeneity beyond this. In particular, no distinction is made between healthy and at-risk members of a given age group. The latter are more likely both to be vaccinated, and to suffer severe outcomes. In other words, our model underestimates the degree to which vaccine is targeted to prevent severe outcomes. Furthermore, at-risk parts of the population often mix associatively with each other rather than the population at large, i.e. in nursing homes and other care facilities that have their own self-contained herd immunity dynamic. This effect is not captured in our model. All of the above limitations contribute to our model somewhat underestimating the effectiveness of vaccination against severe outcomes.

Limitations in stratification notwithstanding, our guiding principle in constructing the model was to include as much of the dynamical “machinery” of influenza as reasonable within a compartmental model. Consequently, the model is readily applicable to other problems beyond the one explored here. In particular, its four-strain structure makes it well-suited to studying scenarios where strain interactions play a key role. An example would be evaluating the impact of a switch from trivalent influenza vaccines (TIVs) to the newly-introduced quadrivalent influenza vaccines (QIVs). For added flexibility, the model allows natural history parameters to be calibrated using one population, then applied to another (broadly similar) population. The simulations we performed here illustrate the usefulness of this capability: A key calibration target, namely the attack rate of influenza in the unvaccinated population (“natural attack rate”) has been calculated in the US but not in Canada. We thus performed our calibration in the former setting, and then use the calibrated inputs in the latter setting.

6. Conclusion

The dynamic, multi-strain model of influenza transmission and vaccination presented here was able to reproduce the impact

on influenza-associated health outcomes of Ontario's transition to universal influenza immunization with reasonable accuracy, while suggesting that some of the assumptions we made regarding influenza vaccine (uptake, duration of protection) may have been conservative. Future studies would benefit from the filling of the corresponding data gaps. The continuing challenge of influenza health burden and the wide array of new vaccines that are becoming available mean that sophisticated, empirically validated dynamic transmission models are needed more than ever before. Such models can assist decision-makers considering the public health and health economic basis for changing recommended vaccines or priority groups.

Conflicts of interest statement

ET and GM are employees of the GlaxoSmithKline group of companies. AC declares that he is a former GSK employee of the GlaxoSmithKline group of companies and is a current employee of Sanofi. AC has stock in both Sanofi and the GlaxoSmithKline group of companies. CB has received consulting fees and grants from GlaxoSmithKline Canada.

Contributors

All authors participated in the design or implementation or analysis, and interpretation of the study: ET, AC, GM, CB contributed to the scientific input, method selection, model development, literature review and review of the study report. ET, CB developed the study report. ET, GM, AC acquired the data. ET, AC, CB performed the statistical analysis. All authors contributed to the development of this manuscript. All authors had full access to the data and gave final approval before submission. All authors agreed to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

Funding source

GlaxoSmithKline Biologicals SA funded this study/research and was involved in all stages of study conduct, including analysis of the data. GlaxoSmithKline Biologicals SA also took in charge all costs associated with the development and publication of this manuscript.

Trademark statement

CANSIM is a registered trademark of Statistics Canada and is an official mark adopted and used by Her Majesty the Queen in Right of Canada as represented by the Minister of Industry Canada.

Acknowledgements

The authors thank Martin Eichner for contributing an idea which was incorporated into the model structure and Heather Santiago (publication manager, GSK) for editorial assistance and manuscript coordination.

Appendix A. Details of the dynamic model

A.1. Age-stratification and demographics

We divided each model compartment into 100 one-year age groups: <1, 1, 2, 3, . . . , 99+ years. This rather high age resolution was chosen for flexibility (some age-stratified model inputs are available with a 1-year age resolution), and also to allow population

aging to be implemented more smoothly. Age-structuring allowed us to implement age-dependent contact patterns within the population in calculating the force of infection (see Appendix A.2), which are specified by means of a contact matrix, see e.g. [6]. Contact matrices can be constructed from survey data; thus far this has been done for many European countries [23] and for the US [24]. The former gives the mean daily number of contacts among members of different age groups, whereas the latter gives the total daily duration of contacts. Either can be utilized in our model, with an appropriate change in the definition of the probability of infection (see Appendix A.2). Denoting the contact matrix as C , element C_{ij} specifies the average number (time) of contacts per day that members of age group i have with members of age group j . Care must be taken in converting contact matrices which utilize age increments larger than the 1 year increment of our model. For example, [24]

	0 – 4 yrs	5 – 9 yrs	...
	C_{11}	C_{21}	...
uses 5-year increments:	5 – 9 yrs	C_{12}	C_{22}
	$\vdots$	$\vdots$	$\ddots$

Assuming that contacts are spread evenly over each age increment, this is converted to 1-year increments as follows:

	0 yrs	1 yrs	2 yrs	3 yrs	4 yrs	5 yrs	...
0 yrs	$\frac{C_{11}}{5}$	...	...	...	$\frac{C_{11}}{5}$	$\frac{C_{21}}{5}$	...
1 yrs	$\vdots$	$\ddots$			$\vdots$	$\vdots$	$\ddots$
2 yrs	$\vdots$		$\ddots$		$\vdots$	$\vdots$	
3 yrs	$\vdots$			$\ddots$	$\vdots$	$\vdots$	
4 yrs	$\frac{C_{11}}{5}$	...			$\frac{C_{11}}{5}$	$\frac{C_{21}}{5}$	...
5 yrs	$\frac{C_{12}}{5}$	...	...	...	$\frac{C_{12}}{5}$	$\frac{C_{22}}{5}$	
$\vdots$	$\vdots$	$\ddots$			$\vdots$	$\ddots$	

Since we intended the model to be usable over an arbitrarily long time horizon, births, all-cause mortality and aging were implemented. Age-specific birth and mortality rates can be specified for the population, and these can be made to vary over time, to reflect historical data or future projections. All demographic processes happen in a discrete step at the end of each flu season:

- (1) All-cause deaths for the season are removed from each compartment.
- (2) The population is aged by one year, by moving the occupants of the age i compartment to the age $i+1$ compartment. The exception is the highest-age compartment, i.e. “99+ yrs”; its occupants remain there until they die.
- (3) Births for the season are added to the lowest-age compartment, i.e. “0 yrs”.

Applying demographic process in a discrete step rather than continuously, while somewhat artificial, has the advantage that it avoids numerical diffusion, which would otherwise produce an unrealistic “smearing out” of age cohorts over time.

A.2. Epidemiology

The model explicitly tracks four distinct types of infection, taken to represent the two dominant A strains (H1N1, H3N2), and the two B lineages (Victoria, Yamagata) of seasonal influenza (denoted in the model as A1, A2, B1 and B2). In keeping with the results of past

studies, we assumed that A and B cross-protection is negligible, and modeled only pairwise cross-protection between the two A strains, and between the two B lineages. The overall structure is thus of a pair of essentially independent two-strain models, each of which we implemented as in [25]. We defined a strain-specific and age-specific susceptibility to infection, $\beta'_{K,i}$: This is the probability that contact with an infected person causes a person in age group i to become infected with strain K , calculated per contact (or per time unit of contact, depending on the definition of the contact matrix; see Section 2.1). If the daily number (respectively time) of contacts of a person in age groups i with members of age group j is C_{ij} , the number of people in age group j infected with strain K is $I_{K,j}$, and the total number of people in age group j is N_j , then the daily force of infection due to strain K for age group i is

$$\lambda_{K,i} = \sum_{j=1}^{100} \beta'_{K,i} C_{ij} \left(\frac{I_{K,j}}{N_j} \right) + \lambda_{\text{ext}} \quad (1)$$

where λ_{ext} is a small additional contribution to the force of infection which represents, in an abstracted way, infection originating from outside the population being modeled. The cause of seasonal variations in influenza incidence continues to be unclear; they may be due to the direct influence of temperature and humidity on transmission, to seasonal changes in contact patterns, or to other factors; see e.g. [11] and references therein. Following the approach of [6], we model this effect as a sinusoidal variation in the force of infection, with a period of one year. This is done by multiplying the force of infection by a factor

$$F_{\text{seasonal}} = \left[1 + A_{\text{seasonal}} \times \cos \left(\frac{2\pi}{365.25} \times (t - t_{\text{max, seasonal}}) \right) \right] \quad (2)$$

which varies between a maximum value of $(1 + A_{\text{seasonal}})$ and a minimum value of $(1 - A_{\text{seasonal}})$ with a period of 1 year (365.25 d), and with the maximum occurring each year at a time $t_{\text{max, seasonal}}$. Both t and $t_{\text{max, seasonal}}$ are in units of days, $t_{\text{max, seasonal}} \in [0, 365.25]$, and $A_{\text{seasonal}} \in [0, 1]$.

Asymptomatic infections were also accounted for in the model. Letting F_{asympt} be the fraction of such infections, the influenza incidence and prevalence output by the model are multiplied by a reduction factor $(1 - F_{\text{asympt}})$ to capture the fact that asymptomatic infections usually remain undetected. The role of asymptomatic individuals in influenza transmission remains a topic of debate; see [26] and references therein. We defined a reduction factor f_{asympt} in the infectiousness of asymptomatic relative to symptomatic individuals and applied this within the model; f_{asympt} can be made to vary from 1 (no reduction) to zero (asymptomatic individuals are incapable of infecting others).

Cross-protection between strains k and l (where k, l are either A1, A2 or B1, B2) is controlled by a pair of parameters $\mathcal{E}_{kl}$ and $\mathcal{E}_{lk}$. $\mathcal{E}_{kl}$ is the probability that immediately after recovering from infection by strain k , an individual is immune not just to strain k but also to strain l . $\mathcal{E}_{lk}$ is the probability of the reverse, i.e. that immediately after recovering from infection by strain l , an individual is immune not just to strain l but also to strain k . Throughout this work, we assumed that cross-protection is symmetric, i.e. $\mathcal{E}_{kl} = \mathcal{E}_{lk}$.

Immunity to influenza strains conferred by infection is only temporary, due to the continual antigenic drift of the influenza virus; see e.g. [10]. We simulated the effect of antigenic drift using an approach common to several other influenza models [6,11,12], whereby individuals lose their immunity to a given strain k at a constant rate θ_k . The inverse, $1/\theta_k \equiv \tau_{k, \text{nat}}$, is then the mean duration of natural immunity.




Fig. 4. The compartments and inter-compartmental flows in “Layer I” of the model, containing all individuals who have not (yet) vaccinated within the current influenza season.

A.3. Vaccination

Vaccination within the model occurs over a finite time window, the yearly start time and length of which are input parameters. Instantaneous or “pulse” vaccination was avoided since this can lead to unrealistic model dynamics when the time of vaccination is not well separated in time from all epidemic peaks. Pulse vaccination can still be approximated if desired (say, for comparison purposes) simply by making the vaccination time window very short.

Both trivalent influenza vaccine (TIV, containing three strains) and quadrivalent influenza vaccine (QIV, containing all four strains) can be simulated within the model. QIV implementation is the most straightforward, since it puts all four strains/lineages on the same footing: The vaccine is simply specified by its efficacy ϵ_k with respect to each $k = A1, A2, B1, B2$. This efficacy can be either a value that remains fixed for the entire simulation, or drawn each flu season from a specified distribution. A person who vaccinates

then becomes “successfully vaccinated” with respect to strain k with probability ϵ_k , and remains susceptible to k with probability $(1 - \epsilon_k)$. In other words, the success of the vaccination is determined at the time of vaccination, rather than the time of challenge. TIV contains only one B lineage, so its B efficacies, ϵ_{B1} and ϵ_{B2} , work differently: Each new season, one of the two B lineages is chosen to be included in the vaccine. The vaccine is then given the lineage-match efficacy with respect to the chosen lineage, and the lineage-mismatch efficacy with respect to the other B lineage.

Since this is a dynamic model, which B strain will actually dominate in a given season is, just as in real life, not known a priori. In real life this prediction is made each season based on worldwide surveillance of influenza strains coordinated by the WHO (e.g. [27]). Mechanistically modeling this selection process is beyond the scope of our work, as it would require a spatial influenza model encompassing the entire planet. Instead, we abstract the process within the model as follows: Each season, we identify the B lineage with the highest prevalence at the exact time that the vaccination




Fig. 5. The compartments and inter-compartmental flows in “Layer II” of the model, containing all individuals who have vaccinated within the current season, and who have not yet had any part of their vaccinated status wane. Note that the compartment $II.V1V2$ is connected only to Layers I and III (see Figs. 7 and 8).




Fig. 6. The compartments and inter-compartmental flows in “Layer II” of the model, containing all individuals who were vaccinated within the current season, but who have had some part of their vaccine-conferred protection wane within that season.

program begins. We define this as the best possible guess of which B lineage to put into the vaccine (which we will denote bBg , “best B guess”). We then define a parameter P_{bBg} , which is the probability with which the “best B guess” is adhered to in a given season. For example, if vaccination commences on October 1, and B1 has higher prevalence than B2 October 1, then B1 will be included in that season’s vaccine with probability P_{bBg} (and B2 with probability $1 - P_{bBg}$). Thus, $P_{bBg} = 0.5$ would correspond to a completely random choice, i.e. “flipping a coin” to pick the B lineage in the vaccine. Since in reality the choice of B lineage must be made more than half a year in advance [27], we expect that realistic accuracy in B lineage prediction will correspond to P_{bBg} well below 1 (the actual value is determined during model calibration; see Appendix B). It is worth noting that even this “best B guess” has an accuracy of well below 100%, since it is very possible for the proportions of B lineages to still change significantly after the vaccination program commences. Indeed, vaccination itself contributes to this, as we explore in Appendix C.

Vaccine-conferred immunity also wanes due to antigenic drift, and this process was implemented in the model analogously to the

waning of natural, infection-conferred immunity, i.e. via a constant rate of immunity loss Θ_k for strain k . This makes $1/\Theta_k \equiv \tau_{k,vacc}$ the mean duration of vaccine-conferred immunity. Since vaccine-conferred immunity is generally thought to be less robust than natural immunity, it was considered important that the model allow the two to wane at different rates.

In a standard compartmental model where vaccination success/failure (assuming efficacy less than 100%) is determined at the time of vaccination, there is no way to distinguish between people who are susceptible because they have not vaccinated yet, and people who remained susceptible after vaccinating unsuccessfully. Thus, if the model applies vaccination over a finite time rather than in an instantaneous “pulse”, there would be no way of preventing a person who unsuccessfully vaccinated (but who would of course have no way of knowing this) from unrealistically vaccinating again within the same flu season. We addressed this by putting anybody who vaccinates – regardless of their success with respect to each of the four influenza strains – into a separate “layer” of the model, which contains all the same compartments connected by (almost) the same differential equations, except that nobody




Fig. 7. The pathways by which vaccination moves individuals from Layer I compartments to Layer II compartments within the model.




Fig. 8. The pathways by which waning of vaccine-conferred immunity (or, in the case of unsuccessful vaccinations, waning of any added protection against outcomes of breakthrough infections; see text) takes individuals from Layer II to Layer III within the model.

in this layer may vaccinate for the rest of the flu season. At the end of the season, all the members of this second layer roll over back into the first layer, thus become eligible for vaccination once again. Another benefit of separately tracking unsuccessful vaccinators is that possible differences in outcomes for “breakthrough” infections (i.e. infections contracted even after vaccination) versus unvaccinated infections can be applied. For example, there is some evidence to suggest that in the case of a break-through infection (or in the terms of our model, of an unsuccessful vaccinator becoming infected), the likelihood of hospitalization is lower than it is for an unvaccinated, infected person, e.g. [28]. Providing the model with the capability to track such an effect necessitated the addition of a third “layer” of compartments to the model; see Appendix A.4 for details.

A.4. Compartment structure and equations

The model considers strain interaction (i.e. cross-protection) between the two A strains and the two B lineages of influenza, but not between influenza A and B, as this is thought to be negligible. Thus, the model actually consists of a pair of independent two-strain components, each having identical structure. The compartments are laid out following the approach of Castillo-Chavez et al. [25], with extra compartments added due to vaccination being considered, as shown in Fig. 4. Since the model is age-stratified, each compartment shown actually comprises 100 individual age sub-compartments. Each compartment is labeled to denote its status with respect to the two influenza strains/lineages. For example:

- V_1V_2 = vaccinated immunity against 1 and 2.
- S_1S_2 = susceptible to 1 and 2.
- I_1V_2 = infected with 1, vaccinated immunity against 2.
- R_1I_2 = natural immunity against (i.e. recovered from) 1, infected with 2.

As mentioned above in Appendix A.3, a proper “book-keeping” of vaccination within the model necessitates the existence of two additional nearly-identical sets of compartments (Figs. 5 and 6).

These three sets of compartments can be thought of as three layers (I, II and III) stacked on top of each other. At the beginning of

a new influenza season, only Layer I is occupied. This continues to be the case until that season’s vaccination program begins. Members of any compartment in Layer I can vaccinate, and subsequently move to Layer II. Possible Layer 1 → 2 pathways are shown in Fig. 7.

As described in Appendix A.3, keeping individuals who have already vaccinated within the current season separate is necessary in order to guarantee that nobody vaccinates more than once a season, which would be unrealistic. This could otherwise arise because, since vaccine efficacy is less than 100%, it is possible to vaccinate with no or only partial success. As an example, a fully-susceptible person (i.e. in compartment S_1S_2) might vaccinate and end up successfully protected against only one of the two strains/lineages (i.e. move to V_1S_2 or S_1V_2), or may be even unluckier and end up with no protection (i.e. remain in S_1S_2). If there were only one set of compartments, there would be no way to distinguish between people who are in, e.g., V_1S_2 after vaccinating in the current season, and those who are left over there from the previous season. The former people would not vaccinate again (since people have no way of knowing if their vaccination was not fully successful), while the latter potentially would. A second set of “already vaccinated this season” compartments, i.e. Layer II, solves this issue, since we do not allow anybody in Layer II to vaccinate.

A related issue arises if we wish to allow for the possibility that even unsuccessful vaccinations confer protection against influenza-associated outcomes, for example, if vaccination reduces the probability of hospitalization in breakthrough infections. This means that we would need to track which of the people who become infected are breakthrough cases (i.e. had previously vaccinated unsuccessfully), so that these people can have a separate outcomes probability (in this example, of hospitalization) applied to them. The above-described 2-layer structure helps here also, but if we consider “unsuccessfully vaccinated” a status that wanes continuously over time, just like successful vaccination, then it is not enough. We have to allow people to leave Layer II due to waning, but we still need to keep them separate from Layer I, so that they do not try to vaccinate again within the current season. We thus added a third layer (Layer III, Fig. 6), which receives all flows due to vaccinated immunity waning from Layer II, and in addition a flow from II. S_1S_2 at the vaccinated-immunity waning rate (to be conservative, we use the higher of the two waning rates, if these differ

between the two strains). In doing this, we made the (conservative) assumption that any benefit conferred by unsuccessful vaccination expires in the time it takes for the vaccinated immunity against a single strain to wane.

At the end of every influenza season, the layers are reset by having all occupants of Layers II and III roll over to the Layer I version of their respective compartment, so that everybody begins the season as “not previously vaccinated this season”.

Below, we present the differential equations which comprise the model. We will do so in an iterative manner, beginning with the equations common to all three layers, then building on these, to mirror the way they are depicted in Figs. 4–8 and implemented in the code. First, the equations for the flows common to all three layers are given in the following equation:

$$\begin{aligned}
 \frac{d(S_1 V_2)}{dt} &= -\lambda_1 F_{\text{seasonal}}(S_1 V_2) + \theta_1 (R_1 V_2) \\
 \frac{d(I_1 V_2)}{dt} &= \lambda_1 F_{\text{seasonal}}(S_1 V_2) - \gamma_1 (I_1 V_2) \\
 \frac{d(R_1 V_2)}{dt} &= (1 - \mathcal{E}_{12}) \gamma_1 (I_1 V_2) \\
 \frac{d(V_1 S_2)}{dt} &= -\lambda_2 F_{\text{seasonal}}(V_1 S_2) + \theta_2 (V_1 R_2) \\
 \frac{d(S_1 S_2)}{dt} &= -(\lambda_1 + \lambda_2) (S_1 S_2) + \theta_1 (R_1 S_2) + \theta_2 (S_1 R_2) \\
 \frac{d(I_1 S_2)}{dt} &= \lambda_1 F_{\text{seasonal}}(S_1 S_2) - \gamma_1 (I_1 S_2) \\
 \frac{d(R_1 S_2)}{dt} &= (1 - \mathcal{E}_{12}) \gamma_1 (I_1 S_2) + \theta_2 (R_1 R_2) - \theta_1 (R_1 S_2) - \lambda_2 F_{\text{seasonal}}(R_1 S_2) \\
 \frac{d(V_1 I_2)}{dt} &= \lambda_2 F_{\text{seasonal}}(V_1 S_2) - \gamma_2 (V_1 I_2) \\
 \frac{d(S_1 I_2)}{dt} &= \lambda_2 F_{\text{seasonal}}(S_1 S_2) - \gamma_2 (S_1 I_2) \\
 \frac{d(R_1 I_2)}{dt} &= \lambda_2 F_{\text{seasonal}}(R_1 S_2) - \gamma_2 (R_1 I_2) \\
 \frac{d(V_1 R_2)}{dt} &= (1 - \mathcal{E}_{21}) \gamma_2 (V_1 I_2) - \theta_1 (V_1 R_2) \\
 \frac{d(S_1 R_2)}{dt} &= (1 - \mathcal{E}_{21}) \gamma_2 (S_2 I_2) + \theta_1 (R_1 R_2) - \lambda_1 F_{\text{seasonal}}(S_1 R_2) - \theta_2 (S_1 R_1) \\
 \frac{d(I_1 R_2)}{dt} &= \lambda_1 F_{\text{seasonal}}(S_1 R_2) - \gamma_1 (I_1 R_2) \\
 \frac{d(R_1 R_2)}{dt} &= \gamma_1 (I_1 R_2) + \gamma_2 (R_1 I_2) + \mathcal{E}_{21} \gamma_2 [(V_1 I_2) + (S_1 I_2)] + \mathcal{E}_{12} \gamma_1 [(I_1 V_2) + (I_1 S_2)] \\
 &\quad - (\theta_1 + \theta_2) (R_1 R_2)
 \end{aligned} \tag{3}$$

The age stratification within the model means that each variable above (except F_{seasonal} is actually a 100-element vector. Thus, for example, the fully written-out form of the first equation is actually

$$\frac{d}{dt} \begin{bmatrix} S_1 V_{2(0 \text{ yrs})} \\ S_1 V_{2(1 \text{ yrs})} \\ \vdots \\ S_1 V_{2(99 \text{ yrs})} \end{bmatrix} = - \begin{bmatrix} \lambda_{1(0 \text{ yrs})} \\ \lambda_{1(1 \text{ yrs})} \\ \vdots \\ \lambda_{1(99 \text{ yrs})} \end{bmatrix} \times \begin{bmatrix} S_1 V_{2(0 \text{ yrs})} \\ S_1 V_{2(1 \text{ yrs})} \\ \vdots \\ S_1 V_{2(99 \text{ yrs})} \end{bmatrix} + \begin{bmatrix} \theta_{1(0 \text{ yrs})} \\ \theta_{1(1 \text{ yrs})} \\ \vdots \\ \theta_{1(99 \text{ yrs})} \end{bmatrix} \times \begin{bmatrix} R_1 V_{2(0 \text{ yrs})} \\ R_1 V_{2(1 \text{ yrs})} \\ \vdots \\ R_1 V_{2(99 \text{ yrs})} \end{bmatrix} \\
 \times \begin{bmatrix} -\lambda_{1(0 \text{ yrs})} S_1 V_{2(0 \text{ yrs})} + \theta_{1(0 \text{ yrs})} R_1 V_{2(0 \text{ yrs})} \\ -\lambda_{1(1 \text{ yrs})} S_1 V_{2(1 \text{ yrs})} + \theta_{1(1 \text{ yrs})} R_1 V_{2(1 \text{ yrs})} \\ \vdots \\ -\lambda_{1(99 \text{ yrs})} S_1 V_{2(99 \text{ yrs})} + \theta_{1(99 \text{ yrs})} R_1 V_{2(99 \text{ yrs})} \end{bmatrix}$$

where denotes an entrywise (or Hadamard–Schur) product. Likewise, all products of vectors throughout these differential equations are entrywise products. The i th element of λ_1 , λ_2 is defined as in Equation 1, with I_{kj} and N_j being the infected and total number of people in the j th age group summed over all 3 layers. The additional external contribution to the force of infection, λ_{ext} , is partitioned between influenza A and B in a way that varies randomly from season to season: $\lambda_{\text{ext},A} = s \lambda_{\text{ext}}$, $\lambda_{\text{ext},B} = (1 - s) \lambda_{\text{ext}}$, with s randomly chosen between 0 and 1. This renders the simulations stochastic. F_{seasonal} is defined as in Eq. (2). As before, θ_1 , θ_2 are the natural immunity waning rates for strains (or lineages) 1 and 2, γ_1, γ_2 are the recovery rates, and $\mathcal{E}_{12}$ and $\mathcal{E}_{21}$ are the cross-protection (we take $\mathcal{E}_{12} = \mathcal{E}_{21}$ throughout). Next, using the set of derivatives above (Eq. (3)) as a starting point, the additional terms in the derivatives due flows within Layer I arising from waning of vaccinated immunity are the following equation:

$$\begin{aligned}
 \frac{d(I \times V_1 V_2)}{dt} &: -(\Theta_1 + \Theta_2) (I \times V_1 V_2) \\
 \frac{d(I \times S_1 V_2)}{dt} &: +\Theta_1 (I \times V_1 V_2) - \Theta_2 (I \times S_1 V_2) \\
 \frac{d(I \times V_1 S_2)}{dt} &: +\Theta_2 (I \times V_1 V_2) - \Theta_1 (I \times V_1 S_2) \\
 \frac{d(I \times S_1 S_2)}{dt} &: +\Theta_1 (I \times S_1 V_2) + \Theta_1 (I \times V_1 S_2) \\
 \frac{d(I \times R_1 V_2)}{dt} &: -\Theta_2 (I \times R_1 V_2) \\
 \frac{d(I \times R_2 S_2)}{dt} &: +\Theta_2 (I \times R_1 V_2) \\
 \frac{d(I \times V_1 R_2)}{dt} &: -\Theta_1 (I \times V_1 R_2) \\
 \frac{d(I \times S_1 R_2)}{dt} &: +\Theta_2 (I \times V_1 R_2)
 \end{aligned} \tag{4}$$

Here, Θ_1 , Θ_2 are the waning rates of vaccine-conferred immunity. Next, again using the set of derivatives in Eq. (3) as a starting point, the additional terms in the derivatives due to waning immunity within Layer III (see Fig. 6) are the following equation:

$$\begin{aligned}
 \frac{d(\text{III} \times S_1 V_2)}{dt} &: -\Theta_2 (\text{III} \times S_1 V_2) \\
 \frac{d(\text{III} \times V_1 S_2)}{dt} &: -\Theta_1 (\text{III} \times V_1 S_2) \\
 \frac{d(\text{III} \times S_1 S_2)}{dt} &: +\Theta_2 (\text{III} \times S_1 V_2) + \Theta_1 (\text{III} \times V_1 S_2) \\
 \frac{d(\text{III} \times R_1 V_2)}{dt} &: -\Theta_2 (\text{III} \times R_1 V_2) \\
 \frac{d(\text{III} \times R_1 S_2)}{dt} &: +\Theta_2 (\text{III} \times R_1 V_2) \\
 \frac{d(\text{III} \times V_1 R_2)}{dt} &: -\Theta_1 (\text{III} \times V_1 R_2) \\
 \frac{d(\text{III} \times S_1 R_2)}{dt} &: +\Theta_1 (\text{III} \times V_1 R_2)
 \end{aligned} \tag{5}$$

Next, the additional terms due to waning flows from Layer II to Layer III (see Fig. 8) are the following equation:

$$\begin{aligned}
 \frac{d(\text{II} \times V_1 V_2)}{dt} &: -(\Theta_1 + \Theta_2)(\text{II} \times V_1 V_2) \\
 \frac{d(\text{III} \times S_1 V_2)}{dt} &: +\Theta_1(\text{II} \times V_1 V_2) \\
 \frac{d(\text{III} \times V_1 S_2)}{dt} &: +\Theta_2(\text{II} \times V_1 V_2) \\
 \frac{d(\text{III} \times S_1 S_2)}{dt} &: -\Theta_2(\text{II} \times S_1 V_2) \\
 \frac{d(\text{III} \times S_1 S_2)}{dt} &: +\Theta_2(\text{II} \times S_1 V_2) \\
 \frac{d(\text{II} \times R_1 V_2)}{dt} &: -\Theta_2(\text{II} \times R_1 V_2) \\
 \frac{d(\text{II} \times R_1 S_2)}{dt} &: -\Theta_1(\text{II} \times R_1 V_2) \\
 \frac{d(\text{II} \times V_1 S_2)}{dt} &: -\Theta_1(\text{II} \times V_1 S_2) \\
 \frac{d(\text{III} \times S_1 S_2)}{dt} &: +\Theta_1(\text{II} \times V_1 S_2) \\
 \frac{d(\text{II} \times S_1 S_2)}{dt} &: -\max(\Theta_1, \Theta_2)(\text{II} \times S_1 S_2) \\
 \frac{d(\text{III} \times S_1 S_2)}{dt} &: +\max(\Theta_1, \Theta_2)(\text{II} \times S_1 S_2) \\
 \frac{d(\text{II} \times R_1 S_2)}{dt} &: -\Theta_2(\text{II} \times R_1 S_2) \\
 \frac{d(\text{III} \times R_1 S_2)}{dt} &: +\Theta_2(\text{II} \times R_1 S_2) \\
 \frac{d(\text{II} \times V_1 R_2)}{dt} &: -\Theta_1(\text{II} \times V_1 R_2) \\
 \frac{d(\text{III} \times S_1 R_2)}{dt} &: -\Theta_1(\text{II} \times V_1 R_2) \\
 \frac{d(\text{II} \times S_1 R_2)}{dt} &: -\Theta_1(\text{II} \times S_1 R_2) \\
 \frac{d(\text{III} \times R_1 R_2)}{dt} &: +\Theta_1(\text{II} \times S_1 R_2) \\
 \frac{d(\text{II} \times R_1 R_2)}{dt} &: -\max(\Theta_1 + \Theta_2)(\text{II} \times R_1 R_2) \\
 \frac{d(\text{III} \times R_1 R_2)}{dt} &: +\max(\Theta_1, \Theta_2)(\text{II} \times R_1 R_2)
 \end{aligned} \tag{6}$$

Finally, the additional terms due to the flows from Layer I to Layer II – that is, all the flows due to vaccination – are given below. To improve readability, given the large number of flows, we have grouped together the terms according to the Layer II compartment they originate from, the same way that they are grouped in Fig. 7:

$$\begin{aligned}
 \frac{d(\text{I} \times V_1 V_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times V_1 V_2) \\
 \frac{d(\text{II} \times V_1 V_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times V_1 V_2) \\
 \frac{d(\text{I} \times S_1 V_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times S_1 V_2) \\
 \frac{d(\text{II} \times V_1 V_2)}{dt} &: +F_{\text{vacc}} \in_1(\text{I} \times S_1 V_2) \\
 \frac{d(\text{II} \times S_1 V_2)}{dt} &: +F_{\text{vacc}}(1 - \in_1)(\text{I} \times S_1 V_2) \\
 \frac{d(\text{I} \times I_1 V_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times I_1 V_2) \\
 \frac{d(\text{II} \times I_1 V_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times I_1 V_2) \\
 \frac{d(\text{I} \times R_1 V_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times R_1 V_2) \\
 \frac{d(\text{II} \times R_1 V_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times R_1 V_2) \\
 \frac{d(\text{I} \times V_1 S_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times V_1 S_2) \\
 \frac{d(\text{II} \times V_1 S_2)}{dt} &: +F_{\text{vacc}} \in_2(\text{I} \times V_1 S_2) \\
 \frac{d(\text{II} \times V_1 S_2)}{dt} &: F_{\text{vacc}}(1 - \in_2)(\text{I} \times V_1 S_2) \\
 \frac{d(\text{I} \times S_1 S_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times S_1 S_2) \\
 \frac{d(\text{II} \times V_1 V_2)}{dt} &: +F_{\text{vacc}} \in_1 \in_2(\text{I} \times S_1 S_2) \\
 \frac{d(\text{II} \times S_1 V_2)}{dt} &: +F_{\text{vacc}}(1 - \in_1) \in_2(\text{I} \times S_1 S_2) \\
 \frac{d(\text{II} \times V_1 S_2)}{dt} &: +F_{\text{vacc}} \in_1(1 - \in_2)(\text{I} \times S_1 S_2) \\
 \frac{d(\text{II} \times S_1 S_2)}{dt} &: +F_{\text{vacc}}(1 - \in_1)(1 - \in_2)(\text{I} \times S_1 S_2) \\
 \frac{d(\text{I} \times I_1 S_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times I_1 S_2) \\
 \frac{d(\text{II} \times I_1 S_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times I_1 S_2) \\
 \frac{d(\text{I} \times R_1 S_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times R_1 S_2) \\
 \frac{d(\text{II} \times R_1 V_2)}{dt} &: +F_{\text{vacc}} \in_2(\text{I} \times R_1 S_2) \\
 \frac{d(\text{II} \times R_1 S_2)}{dt} &: +(1 - \in_2)(\text{I} \times R_1 S_2) \\
 \frac{d(\text{I} \times V_1 I_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times V_1 I_2) \\
 \frac{d(\text{II} \times V_1 I_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times V_1 I_2) \\
 \frac{d(\text{I} \times S_1 I_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times S_1 I_2) \\
 \frac{d(\text{II} \times S_1 I_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times S_1 I_2) \\
 \frac{d(\text{I} \times R_1 I_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times R_1 I_2) \\
 \frac{d(\text{II} \times R_1 I_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times R_1 I_2) \\
 \frac{d(\text{I} \times V_1 R_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times V_1 R_2) \\
 \frac{d(\text{II} \times V_1 R_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times V_1 R_2) \\
 \frac{d(\text{I} \times S_1 R_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times S_1 R_2) \\
 \frac{d(\text{II} \times V_1 R_2)}{dt} &: +F_{\text{vacc}} \in_1(\text{I} \times S_1 R_2) \\
 \frac{d(\text{II} \times S_1 R_2)}{dt} &: +F_{\text{vacc}}(1 - \in_1)(\text{I} \times S_1 R_2) \\
 \frac{d(\text{I} \times I_1 R_2)}{dt} &: -F_{\text{vacc}}(\text{I} \times I_1 R_2) \\
 \frac{d(\text{II} \times I_1 R_2)}{dt} &: +F_{\text{vacc}}(\text{I} \times I_1 R_2)
 \end{aligned} \tag{7}$$

Table 3
Surveillance data for influenza A and B distribution, lineage distribution of influenza B, plus lineage included in TIV ([30], adapted from their Table 5 TIV mismatch seasons are those in which the TIV B lineage matched less than 50% of circulating influenza B (last column; mismatches underlined).

Season	% of A among all influenza viruses	% of B from Yamagata lineage	% of B from Victoria lineage	Lineage included in TIV	% of matched B (= % of included lineage)
1999–2000	99.6%	100%	0%	Yamagata	100%
2000–2001	54%	100%	0%	Yamagata	100%
2001–2002	87%	23%	77%	Yamagata	<u>23%</u>
2002–2003	57%	0.4%	99.6%	Victoria	99.6%
2003–2004	99%	93%	7%	Victoria	<u>7%</u>
2004–2005	75%	74%	26%	Yamagata	74%
2005–2006	81%	22%	78%	Yamagata	<u>22%</u>
2006–2007	79%	24%	77%	Victoria	77%
2007–2008	71%	98%	2%	Victoria	<u>2%</u>
2008–2009	67%	17%	83%	Yamagata	<u>17%</u>
Mean, and bootstrap estimate of 95% CI (1 million resamples)					77% (68%, 86%)
					52% (36%, 82%)

Here, F_{vacc} is the rate at which individuals vaccinate, defined as

$$f_{\text{vacc}} = \begin{cases} -\ln(1 - C_{\text{vacc}}/\Delta t_{\text{vacc}}), & t \in [t_{\text{vacc}}, t_{\text{vacc}} + \Delta t_{\text{vacc}}] \\ 0, & \text{otherwise} \end{cases}$$

where C_{vacc} is the age-stratified vaccine coverage (or uptake), expressed as a fraction of the total population in each age group, Δt_{vacc} is the length of time over which vaccination occurs, and t_{vacc} is the time at which vaccination commences for the current season.

Appendix B. Calibrating the model

B.1. Calibration targets

A fundamental model output, and thus an essential calibration target, is the average yearly fraction of the population infected with influenza, referred to as the attack rate (AR of influenza. We obtained calibration targets for the attack rate in just the unvaccinated part of the population (sometimes called the “natural attack rate”),

$$\text{AR}_{\text{unvacc}} = \frac{\text{number infected while unvaccinated within current season}}{\text{total number unvaccinated within current season}}$$

from a meta-analysis of randomized controlled trials (RCTs) evaluating vaccine efficacy in unvaccinated populations [29]. Relatively speaking, $\text{AR}_{\text{unvacc}}$ is a more universal quantity than the overall AR, since it does not directly depend on the overall population's vaccine uptake. However, due to herd immunity effects, $\text{AR}_{\text{unvacc}}$ still has an indirect dependence on vaccine uptake. Thus, in calibrating the model, it remains important to include a realistic representation of the uptake in the overall population. Most of the RCTs considered by [29] were conducted in various parts of the US; we were unable to identify comparable studies having Canada, or indeed any other country, as the primary source of RCTs. We therefore selected the entire US population as the setting for our calibration. A description of how this setting was implemented within the model is given in Appendix B.2. For each age group's attack rate, the 95% confidence interval obtained by [29] was taken as the allowable target interval.

As a second fitting target, we used the season-by-season partitioning of influenza cases between influenza A and B. This is reported for the US by [30], by season, for the seasons from 1999–2000 to 2008–2009, using data from the CDC's Morbidity and Mortality Weekly Report, and is reproduced in Table 3. We used bootstrapping to estimate the 95% confidence interval for $F(A/A+B)$, the seasonal fraction of all influenza cases which are influenza A averaged over the ten years, then used this interval as our allowable target interval for the same quantity within the model. In using the above data, we made the assumption that the

A – B distribution of the CDC surveillance data is representative of the true A – B distribution, which amounts to assuming that the underreporting fraction is at all times similar for influenza A and B.

As a third fitting target, we used the season-by-season percentage of influenza B matched by the TIV, in other words, the percentage of circulating influenza B which belonged to the lineage included in that season's TIV. These values are also given in Table 3. Again, we used bootstrapping to estimate the 95% confidence interval of the mean, then used that as the allowable target interval. It should be noted that this quantity is not expected to depend on any of the intrinsic properties of influenza in a systematic way; rather, it ought to depend most directly on properties of the vaccination program, specifically the accuracy of the prediction of the upcoming season's dominant B lineage, which is then included in the vaccine. As described in Appendix A.3, the dominant B lineage prediction accuracy in our model is controlled by a parameter P_{bBg} ; we include this as a parameter to be fitted, though strictly speaking it is not a natural history parameter of influenza.

B.2. Performing the calibration

We implemented the setting of the US, from the 2000–2001 to 2008–2009 season, within the model as follows: An initial “burn-in” period of 1990–1999 (during which outputs were not compared to targets) was used. The initial population distribution was that of the entire US in 1990, taken from http://factfinder2.census.gov/faces/tableservices/jsf/pages/productview.xhtml?pid=DEC_00_SF1_QTP1&prodType=table. The population evolved (independent of influenza epidemiology) via birth and death rates given in Tables 4 and 5, linearly interpolated between 1990, 2000 and 2008. US vaccine uptake rates for ages 18 and up were taken from National Health Interview Survey results. In the absence of other data, the uptake rate in ages 2–17 was assumed to be equal to that in the 18–49 yr age

Table 4

US death rate, in deaths per 100,000 per year. Source: <http://www.census.gov/compendia/statab/2012/tables/12s0110.pdf>.

Age (years)	1990	2000	2008
Under 1	1083	807	709
1–4	52	36	32
5–14	29	21	16
15–24	147	115	110
25–34	204	139	142
35–44	310	255	224
45–54	610	543	527
55–64	1553	1231	1105
65–74	3492	2980	2434
75–84	7889	6973	6035
85+	18057	17501	14023

Table 5

US vaccine coverage in percent. Source: http://www.cdc.gov/flu/pdf/professionals/nhis00389_08fluvaxtrendtab.pdf and <http://www.cdc.gov/mmwr/preview/mmwrhtml/rr5908a1.htm> with coverage below 18 years extrapolated using [31].

Age group	2000	2001	2002	2003	2004	2005	2006	2007	2008
6–23 mo	6.0	6.0	6.0	13.0	27.0	26.0	26.0	32.0	32.0
2–49 yr	17.1	15.1	16.3	16.8	17.9	10.4	15.5	17.7	20.0
50–64 yr	34.6	32.1	34.0	36.8	35.9	22.9	33.1	36.2	39.5
65+	64.3	63.0	65.6	65.5	64.6	59.6	64.1	66.4	67.0

group. Vaccine uptake from 6 to 23 months for the 2002–2003 to 2006–2007 seasons was obtained from a MMWR report, <http://www.cdc.gov/mmwr/preview/mmwrhtml/mm5738a2.htm>. We assumed the 2002–2003 rates for all prior seasons, and the 2006–2007 rates for the 2007–2008 and 2008–2009 seasons.

For simplicity, the susceptibility to infection β' was taken to be constant across all ages and influenza strains/lineages. Its value was chosen to yield an approximate seasonal peak basic reproduction number of $R_0 = 1.9$, consistent with a previous estimate for seasonal influenza derived using an age- and strain-structured dynamic model [19]. The relationship between β' and R_0 is as follows:

In an epidemiological model with multiple disease states – in our case, the different age classes $I_1 \dots I_{100}$ – an approximation of the basic reproduction number is given by the dominant (largest absolute value) eigenvalue of the next-generation matrix. The next-generation matrix is defined as follows: Let $F_i(I)$ be the rate at which new infections appear in I_i , and let $V_i(I)$ be the rate at which individuals transfer out of I_i by all other means. Then the next-generation matrix is FV^{-1} , where the matrices F and V have elements given by

$$F_{ij} = \frac{\partial F_i(I)}{\partial I_j}, \quad V_{ij} = \frac{\partial V_i(I)}{\partial I_j},$$

evaluated at the disease-free equilibrium (where the entire population is susceptible). Now for age group i ,

$$\frac{dI_i}{dt} = \lambda_i S_i - (\gamma_i + \mu_i) I_i$$

where μ_i is the mortality rate, as before γ_i is the recovery rate, and from Eq. (1), taking β' as constant across all age groups, as we do in our calibration, setting $N_j = S_j$ (and neglecting λ_{ext}),

$$\lambda_i = \beta' \sum_{j=1}^{100} C_{ij} \left(\frac{I_j}{S_j} \right)$$

To remove the dependence on population distribution, we further take the population in each age group to be equal, i.e. $S_1 = S_2 = \dots = S_{100}$. Then

$$\frac{dI_i}{dt} = \beta' (C_{i1} I_1 + \dots + C_{i100} I_{100}) - (\gamma_i + \mu_i) I_i$$

and so

$$F_{ij} = \beta' C_{ij}$$

and

$$V_{ii} = (\gamma_i + \mu_i), \quad V_{ij} = 0, \quad i \neq j$$

Since we assume a non-age dependent recovery rate $\gamma (= 0.25 \text{ d}^{-1})$, and since the mortality rate is negligibly small compared to the recovery rate, this simplifies to

$$V = \gamma \mathcal{I}$$

Table 6

US Birth rate per 1000 population, used for the US natural history parameter fitting simulations. In using these values in the model, a fixed fraction of 0.4878 of the population is assumed to be female (the 1980 value). Source: <http://www.census.gov/compendia/statab/2012/tables/12s0080.pdf>.

Age of mother (years)	1990	2000	2008
10–14	1.4	0.9	0.6
15–19	59.9	47.7	42.5
20–24	116.5	109.7	106.3
25–29	120.2	113.5	117.5
30–34	80.8	91.2	99.9
35–39	31.7	39.7	46.9
40–44	5.5	8.0	9.8
45–54	0.2	0.5	0.7

where $\mathcal{I}$ is the (100×100) identity matrix, so that $V^{-1} = (1/\gamma) \mathcal{I}$. Thus, in our case,

$$FV^{-1} = \beta' / \gamma C,$$

and so

$$R_0 \approx \rho(FV^{-1}) = \frac{\beta'}{\gamma} \rho(C) = \frac{\beta'}{0.25 \text{ d}^{-1}} 1132.62 = 4530.48 \beta'$$

where the operator ρ denotes the principal eigenvalue, and the contact matrix C is taken from [24]. Thus, we have a linear relationship between β' and an approximate equivalent R_0 .

Table 6 shows all epidemiology- and vaccine-related model input parameters, both those allowed to vary in our calibration, and those held fixed. We performed a set of 5000 fitting runs. These yielded 181 posterior (i.e. accepted) parameter sets. Fig. 9 shows the posterior parameter distributions. The posterior distribution of natural cross-protection is relatively flat for both A and B, meaning that our calibration placed no strong constraints on either H1N1–H3N2 or Victoria–Yamagata natural cross-protection (for the latter we made the prior assumption that vaccine efficacy against the mismatched lineage sets the lower limit for natural cross-protection between the lineages). The posterior distribution of P_{bBg} , the input parameter which controls the accuracy with which each season's dominant B lineage is predicted, suggests that the real-world accuracy of this prediction may actually be higher than currently assumed; see Appendix C.

B.3. A visual check of the model calibration

As a visual check, we plotted a random selection of the 181 posterior (accepted) fitting simulation runs. The World Health Organization (WHO) and US. National Respiratory and Enteric Virus Surveillance System (NREVSS) Collaborating Laboratories have maintained A/B-stratified influenza weekly incidence surveillance starting with the 1997–1998 season (<http://www.cdc.gov/flu/weekly/ussurvdata.htm>). The fitting runs, together with the US surveillance data, are shown in Fig. 1 for the 2000–2001 to 2008–2009 seasons. The surveillance data is greatly scaled down relative to the simulation data, due to incompleteness. Nevertheless, the distribution of seasonal epidemic peaks qualitatively looks similar between simulations and surveillance data. This lends support to the validity of our fitting procedure; that is to say, it appears that the fitting targets are stringent enough to

enforce overall realistic seasonal influenza behavior in the accepted simulations.

Appendix C. Implications for real-world dominant B lineage prediction accuracy

Each season's TIV contains the B lineage predicted (about half a year in advance [37]) to be dominant that season. The chosen lineage has matched the circulating lineage in only about 50% of seasons, from which it has been concluded that the prediction of the dominant B lineage has an accuracy no better than a “coin flip” [27,38]. In our model, a random choice of TIV B lineage amounts to a value of 0.5 for the input parameter which controls this choice, $P_{bBg} \in [0,1]$ (see Appendix A.3 for a definition). In contrast, our model calibration (3) yielded a posterior distribution for P_{bBg} with a mean of 0.69. At the same time, the mean percentage of seasons with a TIV B lineage match among the 181 posterior simulations was 47%, consistent with the US data which had been used as a calibration target (Table 3). This can be understood as follows: Once introduced, the vaccine preferentially suppresses the transmission of whichever B lineage was included in the vaccine; in seasons with significant co-circulation of lineages, this can actually change the dominant B lineage, relative to what it would have been in the absence of vaccination. Thus, in order to yield an *apparent* predictive success rate of 50%, the underlying accuracy of the prediction must in fact be *better* than 50%.

In order to gain further insight into the importance of this effect, we reran the posterior parameter sets yielded by our model calibration from Section 3 a number of times. In the first set of reruns, no vaccination whatsoever was performed; that is, the model simply made a prediction of the dominant B lineage each season, without doing anything to alter the subsequent disease dynamics. We confirmed that $P_{bBg} = 0.5$, i.e. picking a lineage at random, indeed yielded an accuracy of 50%. Next, in order to find the theoretical maximum accuracy of our “best guess”, we used $P_{bBg} = 1$ (i.e. the

best guess is always used), again with no subsequent action (i.e. no vaccination) performed based on the prediction. We obtained an accuracy of 73%.

We then repeated the above, this time turning vaccination on. Now, a random choice of B lineage for the TIV resulted in the chosen lineage matching the ultimately dominant B lineage in only 35% of seasons. Meanwhile, a value of $P_{bBg} = 1$ yielded a match 58% of the time. Comparing these to the above values gives an idea of the large effect vaccination has on which B lineage ultimately dominates a given season. The size of the effect is of course dependent on the vaccine uptake in a given population, as well as on the relative timing of the vaccination program and the epidemic peak. Nevertheless, our results suggest that even under the best possible conditions (recall that $P_{bBg} = 1$ amounts to deferring the choice of TIV B lineage until the day the vaccination program commences every season), the fraction of seasons which will ultimately be recorded as having had a TIV B lineage match will always be significantly below 100%.

Appendix D. Implementing Ontario in the model

D.1. Ontario demographics and vaccine uptake

In order to provide an initial warm-up period, simulations commenced in 1985. The initial population used was thus that of Ontario in 1985, taken from Statistics Canada's CANSIM online database, <http://www5.statcan.gc.ca/cansim/>, Table 051–0001. Birth rates were taken also taken from CANSIM, Table 102–4505; since 2000 was the earliest year available, 2000 birth rates were used from 1985 to 2000. All-cause mortality rates were taken from CANSIM, Table 102–0504.

Ontario vaccine uptake for ages 12 and up was taken from Kwong et al. [2], with the earliest given values, 1996–1997, used all the way from 1985 to 1999 (the next year reported after this is 2000, the first UIIP year). After 2000, uptake rates were




Fig. 9. Posterior distribution of input parameters yielded by the model calibration simulations (5000 prior parameter sets, yielding 181 posterior parameter sets). Each plot's x-axis limits are set equal to the lower and upper bounds of the corresponding prior distribution, all of which were flat (uninformative). Note that the ratio of B/A natural immunity duration is a derived quantity, and was not itself directly varied in the calibration.

Table 7

Age- and time-stratified vaccine uptake in British Columbia, Alberta, Saskatchewan, Manitoba, Quebec and the Atlantic Provinces, from Refs. [2,39].

Year	Age (yrs) 0 (%)	1 (%)	2–11 (%)	12–19 (%)	20–49 (%)	50–64 (%)	65–74 (%)	85+ (%)
1985	1.03	2.05	6.00	6.00	15.00	42.00	54.00	44.00
1999	1.03	2.05	6.00	6.00	15.00	42.00	54.00	44.00
2000	1.03	2.05	9.00	9.00	23.00	58.00	71.00	71.00
2001	1.03	2.05	9.00	9.00	23.00	58.00	71.00	71.00
2003	5.93	11.85	10.00	10.00	29.00	59.00	68.00	71.00
2005	6.55	13.10	14.00	14.00	35.00	62.00	73.00	76.00

linearly interpolated between years for which they were reported (2001, 2003, 2005). Vaccine uptake in the 6–23 month age group in Ontario from 2002 to 2003 onwards was obtained from another source [39]; the 2002–2003 value was very low (1.1% full, 1.9% partial), with a sharp increase thereafter, coincident with the 2003 recommendation from Canada's National Advisory Committee on Immunizations (NACI) to vaccinate all children aged 6–23 months. We used the 2002–2003 uptake for this age group for all earlier seasons, making the assumption that the introduction of the UIIP, coming prior to the NACI recommendation, had negligible effect on the uptake in this age group. Since the model uses age groups one year in size, we applied half the 6–23 month uptake rate to the lowest model age group (<1 yr) to capture the fact that roughly half of this age group was too young to be vaccinated. Also, since [39] gave separate values for fully and partially vaccinated children, we took the effective coverage to be ([% full] + 0.5[% partial]). Uptake in ages 2–11 was assumed to be the same as in the 12–19 age group. All vaccine uptake values used are given in Table 2.

To normalize the Ontario results to the rest of Canada (see Section 4.1), for each simulation, a second version was run using the vaccine uptake in (most of) the rest of Canada during the study period. For ages 12 and up, results were taken from [2]. For ages 2–11, uptake was extrapolated similarly as for Ontario. As we were unable to identify any source for 6–23 month uptake data outside of Ontario, we used the Ontario values. The rest-of-Canada uptake values used are given in Table 7.

D.2. Outcomes probabilities

As explained in Section 4, in the absence of Canadian-specific sources, we made use of US outcomes from [16]. Specifically, these are the probabilities that, as a result of contracting symptomatic influenza, one

- (1) consults a doctor, $P(GP|flu)$;
- (2) is hospitalized, $P(Hosp|flu)$;
- (3) dies, $P(death|flu)$ (i.e. the case fatality ratio).

Additionally, we were interested in the outcome of visiting an emergency room (ER), since this is also examined by [2]. A probability $P(ER|flu)$ is not calculated by [16], so we estimate it for each age group based on the probabilities of the other three outcomes, and the ratio between the rate of ER visits reported by [2], R_{ER} , to the rates of the other three outcomes:

$$P(ER|flu) = \frac{1}{3} \left[P(GP|flu) \times \frac{R_{ER}}{R_{GP}} + P(Hosp|flu) \times \frac{R_{ER}}{R_{Hosp}} + P(death|flu) \times \frac{R_{ER}}{R_{death}} \right]$$

Using non-Canadian outcomes probabilities is not likely to introduce large errors as long as one is only interested – as we are here – in the relative rather than absolute change in outcomes between the pre- and post-UIIP periods. To demonstrate this, we start by considering an age-stratified outcomes probability of a given event,

$$[P_{0y}, P_{1y}, \dots, P_{99y}]$$

(where we have simplified the notation of the probabilities to aid readability), and age-stratified pre- and post-intervention rates of infection

$$[R_{0y}, R_{1y}, \dots, R_{99y}]$$

and

$$[R'_{0y}, R'_{1y}, \dots, R'_{99y}]$$

The pre- and post-intervention age-stratified outcomes rates are then

$$[P_{0y}R_{0y}, P_{0y}R_{1y}, \dots, P_{0y}R_{99y}]$$

and

$$[P_{0y}R'_{0y}, P_{0y}R'_{1y}, \dots, P_{0y}R'_{99y}]$$

respectively. Within a single 1-year age group, the relative rate of events does not depend on the outcomes probability, and is just the ratio of infection rates:

$$\frac{P_i R'_i}{P_i R_i} = \frac{R'_i}{R_i}$$

We can rewrite the age-stratified outcomes probabilities as

$$P_{0y} \times \left[1, \frac{P_{1y}}{P_{0y}}, \dots, \frac{P_{99y}}{P_{0y}} \right]$$

The outcomes rates across all ages are

$$R_{\text{all ages}} = \frac{P_{0y}}{N} \left[N_{0y}R_{0y} + N_{1y}\frac{P_{1y}}{P_{0y}}R_{1y} + \dots + N_{99y}\frac{P_{99y}}{P_{0y}}R_{99y} \right]$$

and

$$R'_{\text{all ages}} = \frac{P_{0y}}{N} \left[N_{0y}R'_{0y} + N_{1y}\frac{P_{1y}}{P_{0y}}R'_{1y} + \dots + N_{99y}\frac{P_{99y}}{P_{0y}}R'_{99y} \right]$$

where $N_{0y} \dots N_{99y}$ are the total number of people in each age group, N is the total number of people summed across all ages, and we have assumed that these numbers are the same pre- and post-intervention. Thus, the RR across all ages is

$$\frac{\left[N_{0y}R'_{0y} + N_{1y}\frac{P_{1y}}{P_{0y}}R'_{1y} + \dots + N_{99y}\frac{P_{99y}}{P_{0y}}R'_{99y} \right]}{\left[N_{0y}R_{0y} + N_{1y}\frac{P_{1y}}{P_{0y}}R_{1y} + \dots + N_{99y}\frac{P_{99y}}{P_{0y}}R_{99y} \right]}$$

The thing to note is that as far as the probabilities are concerned, the RR depends only on the relative variation of the probabilities with age. The same applies to the RR across any subset range of age groups. This is important because it means that any age-stratified outcomes probabilities which differ only by a multiplicative constant will produce the same RRs for the outcome. The relative variation with age of an outcomes probability will tend to be similar across different populations even if the absolute values of that probability are not, especially for two populations with as much in common as Canada and the US.




Fig. 10. Sensitivity analysis: Un-normalized model results (red) for the age-stratified post- versus pre-2000 (=onset of UIIP) relative rates (RRs) of influenza-associated outcomes in Ontario, in a scenario where in ages 2 to 11, the maximum possible change in vaccine uptake occurs, i.e. 0% uptake pre-UIIP, 100% uptake post-UIIP. Filled circles show point estimates, and error bars show 95% CIs. The results obtained by [2] are shown for comparison (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 11. Sensitivity analysis: Un-normalized model results (red) for the age-stratified post- versus pre-2000 (=onset of UIIP) relative rates (RRs) of influenza-associated outcomes in Ontario, in a scenario where for 12 years of age and up the maximum possible change in vaccine uptake occurs, using as limits the 95% CIs for vaccine uptake calculated by [2]. Filled circles show point estimates, and error bars show 95% CIs. The results obtained by [2] are shown for comparison (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 12. Sensitivity analysis: Un-normalized model results (red) for the age-stratified post- versus pre-2000 (=onset of UIIP) relative rates (RRs) of influenza-associated outcomes in Ontario, using the model parameters of Vynnycky et al. [6] and Pitman et al. [7]. Filled circles show point estimates, and error bars show 95% CIs. The results obtained by [2] are shown for comparison (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. Sensitivity analysis: Un-normalized model results (red) for the age-stratified post- versus pre-2000 (=onset of UIIP) relative rates (RRs) of influenza-associated outcomes in Ontario, using the model parameters of Vynnycky et al. [6] and Pitman et al. [7], but with mean duration of vaccinated immunity against influenza A and B reduced to 1 year. Filled circles show point estimates, and error bars show 95% CIs. The results obtained by [2] are shown for comparison (blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

D.3. Sensitivity analysis

One data gap affecting our simulations was the dearth of vaccine uptake data below 12 years of age; as described in [Appendix D.1](#), these values are extrapolated in the model. To explore the theoretical maximum in impact that vaccine uptake changes pre- to post-UIIP could have on outcomes rates in Ontario, we performed a set of simulations as follows: pre-2000, the vaccine uptake from 2 to 11 years of age was set to zero, while post-2000, it was set to 100%. The resulting un-normalized (to allow a more straightforward comparison) RRs are shown in [Fig. 10](#). Comparing to the un-normalized base case shown in [Fig. 2](#), we see a dramatic decrease in the model RRs, many of which now lie below the 95% CI of the results of Kwong et al. [2]. Clearly, our limited knowledge of the below-12 vaccine uptake during the study period has the potential to produce significant inaccuracies in our model results.

Vaccine uptake rates for ages 12 and up do also have a (far smaller) degree of uncertainty; [2] provides 95% CIs for these rates. Similar to above, we considered a scenario of maximal uptake change; this time, however, we had CIs to use as a limit. For the pre-UIIP period, we used the lowest possible rates (i.e. in each age group, the lower bound of the vaccine uptake 95% CI), while for the post-UIIP period, we used the highest possible rate (the upper bound of the 95% CI); the imputed rates below age 12 were kept fixed. Results are shown in [Fig. 11](#). Though RRs are slightly lower than in the baseline set of simulations shown in [Fig. 2](#), the difference is relatively small. Comparably well-constrained uptake rates below age 12 would have substantially firmed up the model results.

For comparison with the three-strain influenza models of Vynnycky et al. [6] and Pitman et al. [7], we also performed a set of simulations using (most of) their baseline natural history parameter values: Mean $R_0 = 1.8$ (as calculated using the next-generation matrix; see [Appendix B.2](#)), mean natural immunity duration of 6 and 12 years for influenza A and B, respectively, vaccinated immunity duration equal to natural immunity duration, no natural cross-protection among strains, and (as also in our baseline simulations) a relative amplitude of ± 0.43 for the seasonal variation in the force of infection. Results are shown in [Fig. 12](#). Compared to the baseline case ([Fig. 2](#)), RRs are substantially reduced. For the RRs taken across all ages, the model point estimates are now all at or below those of Kwong et al. [2]. Note also that the smaller CIs; since here we used fixed values rather than distributions for all natural history parameters, the only variance in outcomes was due to the stochasticity of the simulations.

We performed another set of simulations using the same parameters as above, but reducing the duration of vaccinated immunity back down to one year for both influenza A and B, as in our baseline simulations. The results are shown in [Fig. 13](#). Model RRs are now significantly larger than the results of Kwong et al. We thus see that the longer duration of vaccinated immunity was the main driver of the reduced RRs seen in the previous scenario. In fact, the RRs here also exceed those of our baseline case ([Fig. 2](#)); this is because the higher mean R_0 here (1.8 as compared to the baseline 1.3) reduces the contribution of herd immunity.

References

- [1] Neuzil KM, Hohlbein C, Zhu Y. Illness among schoolchildren during influenza season: effect on school absenteeism, parental absenteeism from work, and secondary illness in families. *Arch Pediatr Adolescent Med* 2002;156(10):986.
- [2] Kwong JC, Stukel TA, Lim J, et al. The effect of universal influenza immunization on mortality and health care use. *PLoS Med* 2008 Oct 28;5(10):e211.
- [3] Longini IM, Halloran ME, Nizam A, Yang Y. Containing pandemic influenza with antiviral agents. *Am J Epidemiol* 2004;159(7):623–33.
- [4] Vynnycky E, Edmunds WJ. Analyses of the 1957 (Asian) influenza pandemic in the United Kingdom and the impact of school closures. *Epidemiol Infect* 2008;136(2):166–79.
- [5] Lambert LC, Fauci AS. Influenza vaccines for the future. *New Engl J Med* 2010;363(21):2036–44.
- [6] Vynnycky E, Pitman R, Siddiqui R, Gay N, Edmunds WJ. Estimating the impact of childhood influenza vaccination programmes in England and Wales. *Vaccine* 2008 Sep 26;26(41):5321–30.
- [7] Pitman RJ, White LJ, Sculpher M. Estimating the clinical impact of introducing paediatric influenza vaccination in England and Wales. *Vaccine* 2012;30(6):1208–24.
- [8] Marjoram P, Molitor J, Plagnol V, Tavar S. Markov chain Monte Carlo without likelihoods. *Proc Nat Acad Sci USA* 2003;100(26):15324–8.
- [9] May R, Anderson RM. Infectious diseases of humans: dynamics and control. Oxford: Oxford University Press; 1991.
- [10] Cox NJ, Bender CA. The molecular epidemiology of influenza viruses. *Semin Virol* 1995 Dec;6(6):359–70.
- [11] Dushoff J, Plotkin JB, Levin SA, Earn DJD. Dynamical resonance can account for seasonality of influenza epidemics. *Proc Nat Acad Sci USA* 2004;101(48):16915–6.
- [12] Finkenstadt BF, Morton A, Rand DA. Modelling antigenic drift in weekly flu incidence. *Stat Med* 2005 Nov 30;24(22):3447–61.
- [13] Van de Velde N, Brisson M, Boily MC. Modeling human papillomavirus vaccine effectiveness: quantifying the impact of parameter uncertainty. *Am J Epidemiol* 2007 Apr 1;165(7):762–75.
- [14] Tully SP, Anonychuk AM, Sanchez DM, Galvani AP, Bauch CT. Time for change? An economic evaluation of integrated cervical screening and HPV immunization programs in Canada. *Vaccine* 2012;30(2):425–35.
- [15] Blower SM, Dowlatabadi H. Sensitivity and uncertainty analysis of complex models of disease transmission: an HIV model, as an example. *Int Stat Rev/Rev Int Stat* 1994;229–43.
- [16] Molinari NA, Ortega-Sanchez IR, Messonnier ML, et al. The annual impact of seasonal influenza in the US: measuring disease burden and costs. *Vaccine* 2007;25(27):5086–96.
- [17] Chowell G, Miller MA, Viboud C. Seasonal influenza in the United States, France, and Australia: transmission and prospects for control. *Epidemiol Infect* 2008;136(6):852–64.
- [18] Nishiura H, Chowell G. The effective reproduction number as a prelude to statistical estimation of time-dependent epidemic trends. In: Chowell G, Hyman J, Bettencourt LS, Castillo-Chavez C, editors. Mathematical and statistical estimation approaches in epidemiology. The Netherlands: Springer; 2009. p. 103–21.
- [19] Truscott J, Fraser C, Hinsley W, et al. Quantifying the transmissibility of human influenza and its seasonal variation in temperate regions. *PLoS Curr* 2009;1.
- [20] Dushoff J, Plotkin JB, Viboud C, et al. Vaccinating to protect a vulnerable subpopulation. *PLoS Med* 2007;4(5):e174.
- [21] Weycker D, Edelsberg J, Elizabeth Halloran M, et al. Population-wide benefits of routine vaccination of children against influenza. *Vaccine* 2005;23(10):1284–93.
- [22] Medlock J, Galvani AP. Optimizing influenza vaccine distribution. *Science* 2009;325(5948):1705–8.
- [23] Mossong JI, Hens N, Jit M, et al. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Med* 2008;5(3):e74.
- [24] Zagheni E, Billari FC, Manfredi P, Melegaro A, Mossong J, Edmunds WJ. Using time-use data to parameterize models for the spread of close-contact infectious diseases. *Am J Epidemiol* 2008;168(9):1082–90.
- [25] Castillo-Chavez C, Hethcote HW, Andreasen V, Levin SA, Liu WM. Epidemiological models with age structure, proportionate mixing, and cross-immunity. *J Math Biol* 1989;27(3):233–58.
- [26] Patrozou E, Mermel LA. Does influenza transmission occur from asymptomatic infection or prior to symptom onset? *Public Health Rep* 2009;124(2):193–6.
- [27] Belshe RB. The need for quadrivalent vaccine against seasonal influenza. *Vaccine* 2010;28 Suppl(4):D45–53.
- [28] Belshe RB, Gruber WC, Mendelman PM, et al. Efficacy of vaccination with live attenuated, cold-adapted, trivalent, intranasal influenza virus vaccine against a variant (A/Sydney) not contained in the vaccine. *J Pediatr* 2000;136(2):168–75.
- [29] Turner D, Wailoo A, Nicholson K, Cooper N, Sutton A, Abrams K. Systematic review and economic decision modelling for the prevention and treatment of influenza A and B. *Health Technol Assess* 2003;7(35):1 (iii–xiii).
- [30] Reed C, Meltzer MI, Finelli L, Fiore A. Public health impact of including two lineages of influenza B in a quadrivalent seasonal influenza vaccine. *Vaccine* 2012;30(11):1993–8.
- [31] Fiore AE, Uyeki TM, Broder K, et al. Prevention and control of influenza with vaccines: recommendations of the Advisory Committee on Immunization Practices (ACIP), 2010. In: MMWR recommendations and reports: morbidity and mortality weekly report recommendations and reports/centers for disease control 59(RR-8); 2010. p. 1.
- [32] Carrat F, Vergu E, Ferguson NM, et al. Time lines of infection and disease in human influenza: a review of volunteer challenge studies. *Am J Epidemiol* 2008 Apr 1;167(7):775–85.
- [33] Tricco AC, Chit A, Soobiah C, et al. Comparing influenza vaccine efficacy against mismatched and matched strains: a systematic review and meta-analysis. *BMC Med* 2013;11(1):153.
- [34] Jefferson T, Di Pietrantonj C, Al-Ansary LA, Ferroni E, Thorning S, Thomas RE. Vaccines for preventing influenza in the elderly. *Cochrane Lib* 2010.
- [35] Jefferson T, Rivetti A, Harnden A, Di Pietrantonj C, Demicheli V. Vaccines for preventing influenza in healthy children. *Cochrane Lib* 2008.

- [36] Jimenez-Jorge S, de Mateo S, Delgado-Sanz C, et al. Effectiveness of influenza vaccine against laboratory-confirmed influenza, in the late 2011–2012 season in Spain, among population targeted for vaccination. *BMC Infect Dis* 2013;13(1):441.
- [37] Palese P. Making better influenza virus vaccines? *Emerg Infect Dis* 2006;12(1):61.
- [38] Ambrose CS, Levin MJ. The rationale for quadrivalent influenza vaccines. *Hum Vaccin Immunother* 2012 Jan;8(1):81–8.
- [39] Campitelli MA, Inoue M, Calzavara AJ, Kwong JC, Guttman A. Low rates of influenza immunization in young children under Ontario's universal influenza immunization program. *Pediatrics* 2012;129(6):e1421–30.