

---

# Evaluating vaccine allocation strategies using simulation-assisted causal modelling

---

Armin Kekić<sup>1</sup> Jonas Dehning<sup>2</sup> Luigi Gresele<sup>1</sup> Julius von Kügelgen<sup>1,3</sup>  
Viola Priesemann<sup>\*,2</sup> Bernhard Schölkopf<sup>\*,1</sup>

<sup>1</sup> Max Planck Institute for Intelligent Systems, Tübingen

<sup>2</sup> Max Planck Institute for Dynamics and Self-Organization, Göttingen

<sup>3</sup> University of Cambridge

\* Equal supervision.

## Abstract

Early on during a pandemic, vaccine availability is limited, requiring prioritisation of different population groups. Evaluating vaccination allocation is therefore a crucial element of pandemics response. In the present work, we develop a model to retrospectively evaluate age-dependent counterfactual vaccine allocation strategies against the COVID-19 pandemic. To estimate the effect of allocation on the expected severe-case incidence, we employ a simulation-assisted causal modelling approach which combines a compartmental infection-dynamics simulation, a coarse-grained, data-driven causal model and literature estimates for immunity waning. We compare Israel’s implemented vaccine allocation strategy in 2021 to counterfactual strategies such as no prioritisation, prioritisation of younger age groups or a strict risk-ranked approach; we find that Israel’s implemented strategy was indeed highly effective. We also study the marginal impact of increasing vaccine uptake for a given age group and find that increasing vaccinations in the elderly is most effective at preventing severe cases, whereas additional vaccinations for middle-aged groups reduce infections most effectively. Due to its modular structure, our model can easily be adapted to study future pandemics. We demonstrate this flexibility by investigating vaccine allocation strategies for a pandemic with characteristics of the Spanish Flu. Our approach thus helps evaluate vaccination strategies under the complex interplay of core epidemic factors, including age-dependent risk profiles, immunity waning, vaccine availability and spreading rates.<sup>1</sup>

## 1 Introduction

The COVID-19 pandemic posed significant challenges to societies and decision makers around the world. Many governments implemented non-pharmaceutical interventions (NPIs) to limit the spread of infections and reduce the number of severe cases [2, 5]. The development of efficient vaccines has provided another key control measure to combat the COVID-19 pandemic [21]. However, vaccine supply can fail to meet demand, and vaccine uptake can fall short of expectations. Under these conditions, governments have to find rational strategies to allocate vaccines to minimise harm. It is therefore important to understand how to evaluate and compare different vaccine allocation strategies.

Given an observed evolution of infections and severe cases, we seek to answer central policy questions: *Given limited vaccine availability, should one have prioritised first vaccinations of the*

---

<sup>1</sup>Link to source code: <https://github.com/akekic/covid-vaccine-evaluation>

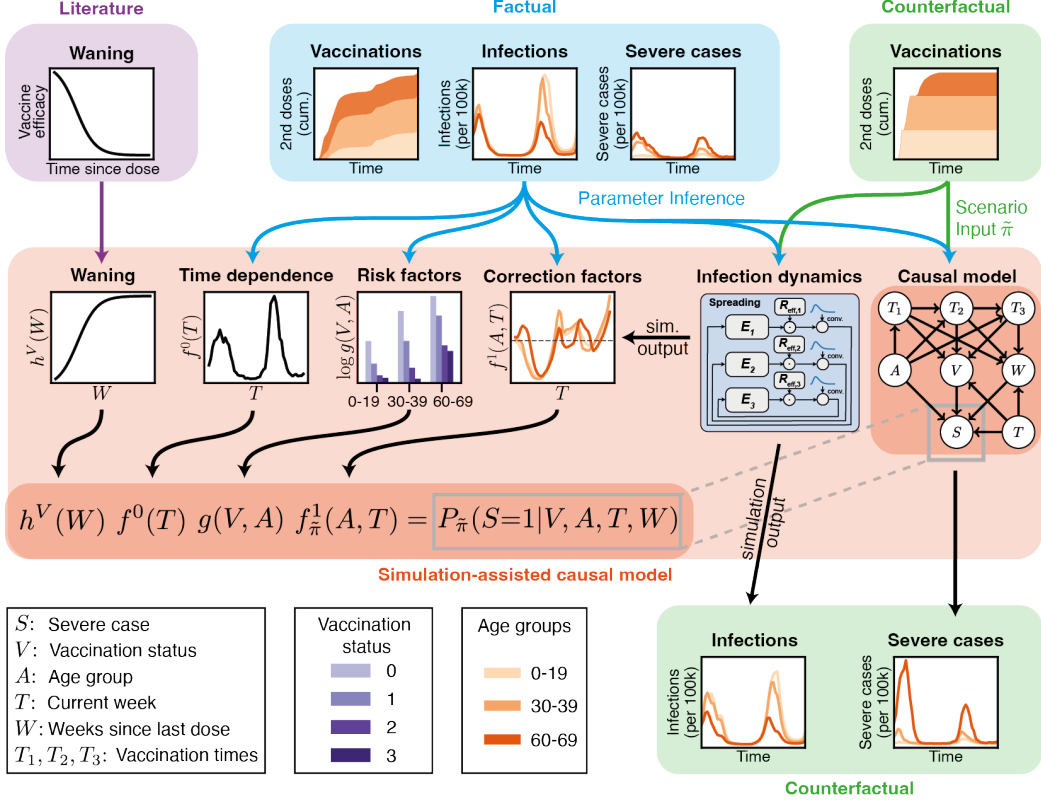


Figure 1: **Method overview.** Our goal is to compute the weekly severe-case incidence (bottom right) under counterfactual vaccine allocation strategies (in the example shown here: *YoungFirst*, top right). To compute this counterfactual scenario, we provide an estimate of the severity mechanism  $P_{\bar{\pi}}(S=1|V, A, T, W)$  for all combinations of  $(V, A, T, W)$  through the proposed factorisation (2). The other conditionals in the causal graph are directly estimated from data (factual strategy, top centre) or are intervened upon according to the counterfactual strategy. The risk factors  $g(V, A)$  and the time dependence  $f^0(T)$  are estimated from data after accounting for immunity waning  $h^V(W)$  derived from literature estimates (top left). An SEIR-like infection dynamics model is fit to the factual infections and subsequently used to simulate infections under the counterfactual strategy. The simulation output is used to compute the correction factors  $f^1_{\bar{\pi}}(A, T)$  accounting for the age-specific change in probability of being infected. Only three of nine age groups are shown for simplicity.

*young or booster shots for the elderly? Which age groups should have been targeted preferentially to increase vaccine uptake? How would a different age-dependent risk profile have impacted the outcome of a vaccination policy?* Answering such questions requires computing the effects of hypothetical interventions on a complex dynamical system (given observations of the same system under different conditions). Such *What if?*-scenarios lie at the heart of causal inference and relate to *counterfactual* reasoning: that is, reasoning about how events would have turned out had some circumstances been altered, all else being equal [22, 26].

The gold standard for inferring *average* causal effects are randomised controlled trials (RCTs) [8]. However, to compare full vaccine allocation strategies at the country level, running an RCT is infeasible in practice (there is only one copy of each country), mirroring similar challenges in the empirical estimation of *individualised* treatment effects [16, 27]. Moreover, it would be ethically unacceptable to implement vaccine allocation strategies that are expected to be suboptimal for the population. Here, modelling approaches provide an important tool to fill this gap.

To capture the effects of changes in vaccine allocation strategy, we have to model their impact on spread and hospitalisation. Furthermore, we need to consider aspects such as vaccine efficacy, immunity waning, age-dependent risk profiles and contact structures. Two established modelling paradigms are compartmental differential equation models and machine learning (ML) approaches.

In principle, compartmental models like the susceptible-infectious-recovered (SIR) model (and its extensions) can be used to answer the types of questions we are interested in—provided that all relevant parameters are known sufficiently well [29, 31, 10, 15, 1]. However, this is typically not the case. Thus, a framework combining inference of parameters and prediction is necessary. To jointly model infections and severe cases, compartmental models require a large state space, whose parameters can be difficult to estimate from data without overfitting [24].

ML methods excel at fitting data and making predictions based on statistical associations, but are generally unable to answer causal questions. Moreover, they are unreliable when the underlying data distribution changes. Yet, we are precisely interested in how our system behaves under distribution shifts: we want to know how the expected severe-case incidence would have changed had we implemented different vaccine allocation strategies.

Causal models occupy a middle ground between the two paradigms and are better-suited for our purpose since they are *modular* and *interventional*. Modularity refers to a model being composed of autonomous components or subsystems and allows for combining different sources of knowledge. Interventions are naturally supported by causal models since they explicitly capture the data-generating mechanisms rather than mere statistical associations. Causal models can thus answer counterfactual questions while clearly expressing the underlying assumptions [22, 23]. However, using them to describe time-varying systems and modelling the dynamics of epidemic spread may be cumbersome.

Since neither compartmental nor ML nor causal models on their own are fully suited for our task, we resort to a hybrid modelling approach involving a modular combination thereof. We propose a coarse-grained causal model in which most components are estimated from data, but where additionally one of the modules is a compartmental model and another one is derived from literature estimates. In particular, we use a susceptible-exposed-infectious-recovered (SEIR) model for the infection dynamics—but not for severe cases—and rely on literature estimates of immunity waning, as illustrated in Fig. 1. The SEIR-like infection-dynamics model can be fit more easily to data compared to a joint model of severe cases and infections. For severe cases, given a set of qualitative causal assumptions, fitting the causal model reduces to the problem of statistical estimation of conditional probabilities. This combines the strengths of the two approaches: the data-driven nature of causal models and the expressivity of compartmental models.

As a case study, we apply our method to a comprehensive dataset collected in Israel [19]. Specifically, we compare several counterfactual age-dependent vaccine allocation strategies to the factual strategy, assuming a fixed number of administered doses and fixed vaccine uptake rate per age group. We also simulate the effect of campaigning for vaccine uptake in a given age group by increasing vaccine uptake rate in one group and computing the effect on the severe-case incidence across all age groups. To showcase the capability of our model to change and examine the influence of individual modules, we consider a different type of disease whose age-dependent risk profile is based on the Spanish Flu.

## 2 Simulation-assisted causal modelling

### 2.1 Causal model

We use a causal graphical model [22], as shown in Fig. 1 (middle right), to describe an individual’s probability of developing severe COVID-19  $S \in \{0, 1\}$  where 1 indicates a severe case and 0 indicates a mild case or no infection at all. We assume the severe-case probability depends on the following variables: the vaccination status  $V \in \{0, 1, 2, 3\}$  indicating the number of vaccine doses a person has received; their age group  $A \in \{0-19, 20-29, \dots, 80-89, 90+\}$ ; the current week  $T \in \{1, \dots, M\}$ ; and the time since the last dose was received  $W \in \{1, \dots, M\}$ . The variables  $V$  and  $W$  are functions of the weeks in which the respective doses were received  $T_i \in \{1, \dots, M+1\}$  for  $i \in \{1, 2, 3\}$  and the current week  $T$ .

We want to estimate the causal effect of changing the distribution of vaccination times for different age groups  $A$  on the expected severe-case probability:

$$\mathbb{E}[S | \text{do}(T_1, T_2, T_3 \sim \tilde{P}(T_1, T_2, T_3 | A))] \quad (1)$$

where we denote the distributions of vaccination times pre- and post-intervention as  $\pi = P(T_1, T_2, T_3 | A)$  and  $\tilde{\pi} = \tilde{P}(T_1, T_2, T_3 | A)$ , respectively.

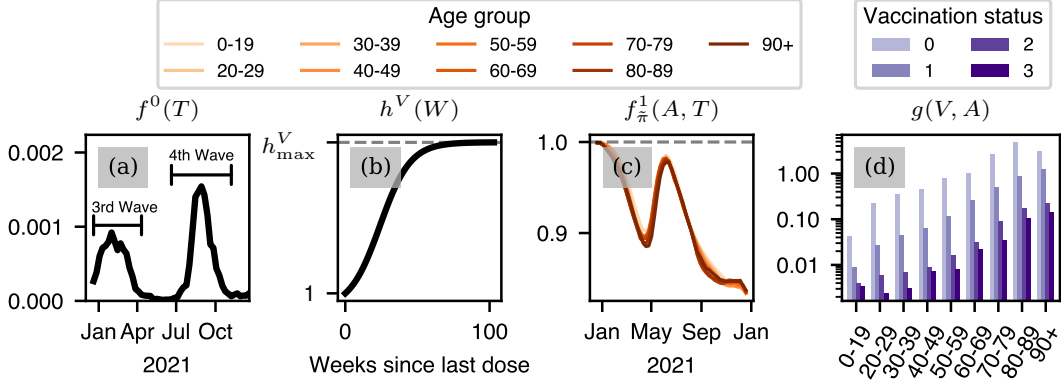


Figure 2: **Factors contributing to the severe-case probability.** (a) Estimated time dependence  $f^0(T)$  approximately following the two infection waves in Israel in 2021. (b) The waning curve  $h^V(W)$  with increasing risk over time as the immunity wanes. The waning curve is computed based on results from [30]. (c) Estimated age-dependent correction factors  $f_{\pi}^1(A, T)$  accounting for the change in population-level infection dynamics. The correction factors shown here correspond to the scenario with increased vaccine uptake rate by 2%, leading to a relative decrease in the number of infections. (d) Risk factor  $g(V, A)$  estimates indicating the relative risk of having a severe case by age  $A$  and vaccination status  $V$ .

There are three challenges we have to address: (i) In the publicly available data by Israel’s Ministry of Health [19] severe outcomes are not registered as a function of time since the last dose was received. We only have access to the marginal distribution  $P(S=1|V, A, T) = \sum_w P(S=1|V, A, T, w)P(w|V, A, T)$ . (ii) Computing the post-intervention severity (1) can involve evaluating the conditional  $P(S=1|V, A, T)$  for combinations  $(V, A, T)$  for which there are no observations. In particular, due to the implemented age-ranked vaccine allocation strategy in Israel, there may not have been any vaccinated subjects in certain younger age groups for some of the early weeks. (iii) Changing the vaccine allocation strategy influences the probability of having a severe case in two ways: firstly, by changing the probability of having immunity through vaccination and secondly, through impacting the infection dynamics on the population level. Such changes in infection dynamics are not captured by the causal model alone.

To address these challenges, we propose a factorisation of the *severity mechanism*

$$P_{\pi}(S=1|V, A, T, W) = f^0(T) g(V, A) h^V(W) f_{\pi}^1(A, T). \quad (2)$$

The waning of the immunity against infection is described by  $h^V(W)$ ; it can be derived from literature estimates for the vaccine efficacy against infection as a function of time since the last dose was administered [30]. The observed aggregate time dependence of the probability of having a severe case is captured by  $f^0(T)$ . The factor  $g(V, A)$  describes the age- and vaccination-status dependent relative risk factor of having a severe case.  $f^0(T)$  and  $g(V, A)$  can be estimated from data after correcting for the influence of waning.  $f_{\pi}^1(A, T)$  is a correction factor that depends on the post-intervention vaccination distribution and accounts for the change in infection dynamics. The correction factor couples the causal model and the infection dynamics simulation.

The factorisation (2) resolves the challenges above by allowing us to (i) incorporate literature knowledge about immunity waning into our causal model, (ii) estimate  $P(S=1|V, A, T, W)$  for all values in the conditioning set by transferring knowledge between vaccination states, age groups and weeks, and (iii) take into account the population level impact on the infection dynamics.

## 2.2 Estimating the severity mechanism factors

**Risk factors for the unvaccinated.** Estimating the factorisation (2) is ill-posed because the overall scale of the factors  $g$ ,  $h^V$  and  $f^0$  is not well-defined.<sup>1</sup> We remove this ambiguity by setting  $g(V=0, a^*)=1$  for the age group  $a^*=60-69$ . Note that there is no waning for  $V=0$ , hence  $P(S=1|V=0, A, T, W) = P(S=1|V=0, A, T)$ . We can then estimate the other risk factors of the unvaccinated by

$$\hat{g}(0, a) = \frac{\mathbb{E}_T[P(S=1|V=0, A=a, T)]}{\mathbb{E}_T[P(S=1|V=0, A=a^*, T)]}. \quad (3)$$

<sup>1</sup>We can double  $f^0$  and halve  $g$  and end up with the same value for the severity mechanism.

**Immunity waning curve.** The vaccine efficacy against infection as a function of time since the administration of the second dose of the BionTech vaccine is reported in [30]. Let  $\text{VE}^v(w)$  be the vaccine efficacy  $w$  weeks after receiving the  $v^{\text{th}}$  dose. The waning curve  $h^v$  is given by (full derivation is shown in App. A)

$$h^v(w) = \frac{1 - \text{VE}^v(w)}{1 - \text{VE}^v(0)}. \quad (4)$$

**Risk factors for the vaccinated.** In the data [19] we can only observe  $P(S=1|V, A, T)$  since we do not have data on severe outcomes as a function of the time since the last dose  $W$ . However, we do have data on the distribution of times since the last dose was received  $P(W|V, A, T)$ . This allows us to use the waning function (4) to estimate the risk factors for  $V=v > 0$  since

$$P(S=1|v, a, t) = \sum_w P(S=1|v, a, t, w)P(w|v, a, t) = f(t)g(v, a) \sum_w h^v(w)P(w|v, a, t). \quad (5)$$

This motivates the following estimator:

$$\hat{g}(v, a) = \frac{\mathbb{E}_T [P(S=1|V=v, A=a, T) \mathbb{E}_{W|V=v, A=a, T} [h^v(W)]^{-1}]}{\mathbb{E}_T [P(S=1|V=0, A=a, T)]} \hat{g}(0, a), \quad (6)$$

*i.e.*, we correct for the waning that occurred in the population to estimate the risk factor under full immunity.

**Time dependence.** After correcting for immunity waning and differences in risk factors, we can estimate the overall time dependence:

$$\hat{f}(t) = \mathbb{E}_{V, A|T=t} \left[ \frac{P(S=1|V, A, T=t)}{\hat{g}(V, A) \mathbb{E}_{W|V, A, T=t} [h^V(W)]} \right]. \quad (7)$$

**Infection dynamics correction factor.** It can be shown (see App. B) that the correction factor  $f_{\tilde{\pi}}^1(A, T)$  under the counterfactual vaccine allocation strategy is given by the relative change in weekly infection probability for each age group

$$f_{\tilde{\pi}}^1(A, T) = \frac{P_{\tilde{\pi}}(I=1|V, A, T, W)}{P_{\pi}(I=1|V, A, T, W)}. \quad (8)$$

Estimated factors are shown in Fig. 2.

### 2.3 Infection dynamics simulation

To estimate the effect of changing the vaccine allocation strategy on the infection dynamics ( $P_{\tilde{\pi}}(I=1|V, A, T, W)$ ) we first infer the parameters of a Bayesian SEIR-like model under the factual strategy  $\pi$ . For each age group  $a$ , we fit a time-dependent *base reproduction number* (Fig. 3 left), *i.e.* the reproduction number in a hypothetical non-immune population. From this base reproduction number, we calculate the *effective reproduction number*  $R_{\text{eff}, a}(t_{\text{day}})$  which includes the vaccination effect. With this we can model the infection dynamics with a *discrete renewal process* [11, 9, 2]:

$$E_a(t_{\text{day}}) = \sum_{a'=1}^9 \sqrt{R_{\text{eff}, a}(t_{\text{day}})} C_{a, a'} \sqrt{R_{\text{eff}, a'}(t_{\text{day}})} \sum_{\tau=0}^{10} E_{a'}(t_{\text{day}} - 1 - \tau) \tilde{g}(\tau), + h_a(t_{\text{day}}), \quad (9)$$

$$S_a(t_{\text{day}}) = S_a(t_{\text{day}} - 1) - E_a(t_{\text{day}} - 1), \quad (10)$$

$$\tilde{g}(\tau) = \text{Gamma}(\tau; \mu=4, \sigma=1.5). \quad (11)$$

Here  $E_a(t_{\text{day}})$  is the number of newly exposed individuals on day  $t_{\text{day}}$ , who will later become infectious; It is implicitly modelled by the latent kernel  $\tilde{g}$ . The infections from one age group to another are encoded by a contact matrix  $\mathbf{C}$ .  $S_a(t_{\text{day}})$  is the number of susceptible people and  $h_a(t_{\text{day}})$  is an external influx due to travellers entering Israel.

By delaying the number of exposed people, we estimate the number of cases, and define a likelihood function by comparing them with the observed cases. By defining all the necessary prior distribution (see App. C) we can infer the parameters of the model including the time-dependent reproduction numbers in Bayesian manner with Monte-Carlo methods.

With the inferred reproduction number we rerun the model with a counterfactual strategy  $\tilde{\pi}$  to obtain an estimate of the number of infections in the counterfactual scenario. The ratio of the number of infections under the factual and counterfactual scenario then defines the correction factor (eq. 8).

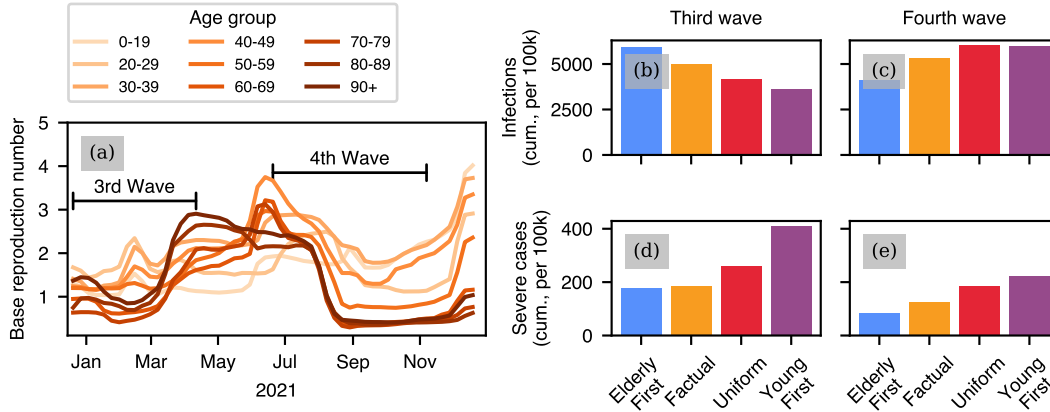


Figure 3: Plot (a) shows the base reproduction numbers inferred by the SEIR-like infection dynamics model. The base reproduction numbers describe the contribution of each age group to infection spread after accounting for the effect of vaccination. Bar plots show the cumulative incidences of infections (b, c) and severe cases (d, e) for the two infection waves in 2021 under the factual and counterfactual vaccine allocation strategies. For the third and fourth wave we sum all cases from 2020-12-20 to 2021-04-11 and 2021-06-20 to 2021-11-07, respectively.

### 3 Results

#### 3.1 Counterfactual vaccine allocation strategies

We compare four age-dependent vaccine allocation strategies: the factual and three counterfactual ones. (i) **Factual**: The vaccine allocation strategy implemented in Israel generally prioritised the elderly (starting with all aged 60 and over), but also prioritised nursing home residents, patients with pre-existing medical conditions, and front-line health care workers [25]. (ii) **Uniform**: All age groups are vaccinated at the same rate. (iii) **ElderlyFirst**: We prioritise age groups in descending order, starting with the oldest. This strategy differs from the factual strategy in that it strictly prioritises by age and does not consider other factors. (iv) **YoungFirst**: The opposite of ElderlyFirst.

**Infections.** During the third wave, YoungFirst leads to the lowest infection incidence (Fig. 3 (b)). The most effective strategy in preventing infections is determined by which age group is contributing most to infection spread. As shown in Fig. 3(a), during the third wave, the base reproduction numbers tended to be higher among the young and middle age groups and lower for the elderly, pointing to differences in behaviour and average number of contacts. Therefore, strategies that prioritise young and middle age groups are most effective at preventing infections. In the fourth wave, the relative effectiveness of strategies at preventing infections is reversed and ElderlyFirst leads to the lowest infection incidence (Fig. 3 (c)). This is due to changes in the base reproduction numbers per age group: For the first half of the wave, the base reproduction numbers in the youngest age group, accounting for roughly a third of the population, tend to be lower than for the other age groups due to the school summer break. Conversely, the base reproduction numbers for the middle-aged and elderly tend to be higher. Therefore, prioritising these age groups is a more effective measure for preventing infection spread at that point in time.

**Severe cases.** The ElderlyFirst strategy leads to the lowest cumulative *severe* case incidence; slightly lower than for the Factual strategy for both waves (Fig. 3 (d,e)). The Uniform and YoungFirst strategies lead to the highest severe-case incidence. The severe-case risk is lowest for the youngest age groups (Fig. 2(d)). Hence, prioritising the younger age groups leaves the older age groups (which are at higher risk) less protected against severe cases. While the YoungFirst strategy reduces the severe-case incidence in the youngest age group, this reduction is far outweighed by the increase in the other age groups. We find the opposite effect for the ElderlyFirst scenario: we have a reduction of the severe-case incidence in the elderly and an increase in the younger age groups.

Our results suggest that under the ElderlyFirst strategy, the cumulative severe-case incidence could have been reduced had this strategy been implemented in Israel (177 vs. 184 and 84 vs. 126 per 100k in the third and fourth waves, respectively). However, our counterfactual vaccine allocation strategies

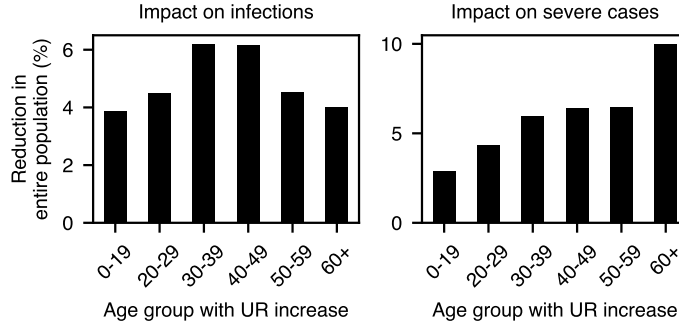


Figure 4: **Impact of increasing vaccine uptake rate (UR) in a given age group.** In each scenario, the vaccine uptake rate is increased in a given age group by a fixed number corresponding to 0.6% of the population ( $N=55,746$  doses) being motivated to get vaccinated. We assume that the change comes from originally unvaccinated individuals who are persuaded to receive three doses. All other age groups follow the factual vaccine allocation strategy. The additional doses are spread over the entire time period by scaling the weekly administered doses by a constant factor. The plots show the impact on cumulative infections (left) and severe cases (right) in the entire population—not just in the age group in which the UR was increased. We consider cases from 2020-12-20 to 2021-12-25.

make some simplifying assumptions that are difficult to implement in practice. We assume that it is possible to vaccinate all willing patients of an age group before moving on to the next age group without delays. In practice, this is difficult to accomplish, given how hard it can be to reach the elderly in particular. Therefore, no realistic vaccine allocation strategy can be as strict as the protocol followed in *ElderlyFirst* and will have some overlap between age groups. Hence, from our counterfactual results and these observations, we conclude that the *Factual* strategy may have been close to optimal.

### 3.2 Impact of increasing vaccine uptake

We simulate the effect of increasing the vaccine uptake in a single age group. Limited vaccine uptake is a factor that prolongs the necessity of non-pharmaceutical interventions [21]. One possible intervention governments have at their disposal is to encourage vaccinations through advertisement campaigns, which can be targeted at specific age groups. Our method allows us to estimate the impact of increasing vaccine uptake on the expected incidence of severe outcomes.

Increasing vaccine uptake in the oldest age groups is most effective in decreasing the severe-case incidence in the total population (Fig. 4 (right)). On the other hand, increasing vaccine uptake in the middle age groups is most effective in decreasing the infection incidence in the total population. There are two effects of an increase in vaccine uptake. Firstly, there is a larger number of individuals who are better protected against infections and severe cases through vaccine-induced immunity. This effect can be explained by the risk factors shown in Fig. 2(d): by increasing the vaccine uptake rate in a given age group, we effectively move an additional part of this population from vaccination status 0 (unvaccinated) to 3 (boostered). We find the largest effect for the elderly, since they have the largest absolute difference in risk factors between vaccination states 0 and 3. Secondly, there is a population level effect: by influencing the infection dynamics, the total number of infections is reduced (Fig. 4 (left)). Since the base reproduction number tends to be higher for the middle age groups (Fig. 3 (a)) they have a larger impact on the number of infections when their vaccine uptake rate is increased. However, when considering the severe-case incidence, this effect on infections is not large enough to outweigh the differences in risk factors (Fig. 4 (right)). In summary, even when taking the effect on infection dynamics into account, it would have been most beneficial to increase vaccine uptake rate in the elderly.

### 3.3 Simulating other disease types

To investigate if one can generalise the recommendation to first vaccinate age groups with the highest severe-case risk, we simulate different types of diseases. The explicit factorisation of the severity mechanism (2) allows us to dissect the different contributing factors that determine the expected severe-case incidence. One of those factors is the age- and vaccination-dependent risk profile of Covid-19. We can ask: *How would a disease with a different risk profile behave?* We compare three disease types (Fig. 5 (left)): (i) COVID-19, (ii) Spanish Flu: We use age-specific excess




Figure 5: **Comparison of disease types.** The left column (a-c) shows the age-specific risk factors for the unvaccinated for three considered disease types. The vaccine efficacy is assumed to be independent of age and the same across all disease types. The risk factors are normalised across disease types to lead to the same number of severe cases under the Uniform strategy. The other plots show the cumulative incidences of infections (d) and severe cases (e) under different vaccine allocation strategies between 2020-12-20 and 2021-04-11 (third wave).

respiratory death rates associated with the Spanish Flu pandemic in Kentucky 1918-1919 [32] as an approximation for the risk factors  $g(0, A)$  for the Spanish Flu. Age groups 20-29 and 30-39 experience the highest severe-case risk. To obtain the other risk factors  $g(V > 0, A)$ , we assume a constant vaccine efficacy for all age groups. The other factors of the severity factorisation (2) are assumed to be the same as for COVID-19 (shown in Fig. 2). (iii) Flat Risk<sup>2</sup>: All age groups have the same severe-case risk.

Besides the vaccine allocation strategies shown in Sec. 3.1 we consider two additional strategies that take the altered risk profiles into account: (i) RiskRanked: We prioritise age groups in descending order of the risk factors. (ii) RiskRankedReversed: The opposite of RiskRanked.

There are no differences between diseases regarding the cumulative *infection* incidence under the four initial vaccine allocation strategies, *i.e.* Factual, Uniform, ElderlyFirst and YoungFirst. Of those, the strategies that prioritise young and middle age groups are most effective at preventing infection spread during the third wave, as discussed in Sec. 3.1. For COVID-19, the RiskRanked and RiskRankedReversed strategies are most similar to the Factual and YoungFirst strategies, respectively, and lead to similar infection incidences. For the Flat Risk, RiskRanked and RiskRankedReversed are identical to the Uniform strategy and lead to the same infection outcome. For the Spanish Flu, the most effective strategy for preventing infection spread is RiskRanked as it prioritises age groups 20-29 and 30-39 which tend to have the highest base reproduction numbers during the third wave (see Fig. 3(a)); Conversely, RiskRankedReversed is the least effective.

For COVID-19, the severe-case incidence is lowest for strategies that prioritise high-risk age groups and lowest for strategies that do the reverse, as discussed in Sec. 3.1. The same can be said for the Spanish Flu, however, the adversarial strategy RiskRankedReversed leads to even worse outcomes. For COVID-19, younger and middle age groups have high base reproduction numbers during the third wave and the elderly had high severe-case risk. For the Spanish Flu, however, the age groups with high severe-case risk and high base reproduction numbers are the same (20-29 and 30-39). Therefore, the RiskRankedReversed strategy is adversarial in two ways: it both leads to a high infection and severe-case incidence.

<sup>2</sup>The factual vaccine uptake rate is influenced by the age-specific risk structure of COVID-19. Since the elderly are at higher risk of severe case and death, they have more incentive to get vaccinated. To remove this bias from the simulation set-up, we assume a flat vaccine uptake rate of 90% willingness to receive all three vaccine doses across all age groups for the Flat and Spanish Flu risk profiles.



In summary, this simulation shows that, even when taking infection dynamics into account, following a strategy in which people most at risk are vaccinated first leads to the lowest severe-case incidence.

## 4 Discussion and related work

Previous work that retrospectively evaluated vaccination campaigns focused mainly on the overall success by modelling counterfactual scenarios with fewer or no vaccinations [33, 14]. In addition, there have been simulation studies conducted before the start of the vaccination campaign that, similar to our work, evaluate different age-dependent vaccine allocation strategies [3, 18, 10]. The forward-looking studies find that in general, the middle age groups have the biggest impact on *infection* incidence (*e.g.* Fig. 1 in [3], for demographics corresponding to the United States), which is in agreement with our findings (see Fig. 4 (left)). However, in these studies, such differences in contribution to spread stem from assumptions on the contact matrix [3, 10] or prior knowledge of susceptibility [3, 18, 10], whereas in our method, we estimate the age-dependent reproduction numbers directly from infection data. The agreement suggests that both approaches lead to qualitatively similar results. Other studies investigating age-dependent transmission inferred from observed cases [4] or seropositivity data [28] also find that middle age groups have the largest impact on transmission. In the present work, by retrospectively taking observed infection data into account, we were additionally able to show that the timing of the vaccination campaign relative to non-pharmaceutical interventions is crucial. In the third wave for instance, with partially open schools in Israel, vaccinating the young would have been most effective at preventing infections; whereas, in the fourth wave, with school holidays during the start of the wave, this strategy would have been among the least effective (see Fig. 3). Nevertheless, when it comes to minimising *severe cases*, all studies agree that (under realistic parameter settings [18]) prioritising the elderly is most effective due to the large difference in the infection-fatality ratio [3, 18, 10] (see Figs. 3 and 4 (right)).

While we were able to include many factors relevant to severe cases and infections, our method has some limitations. Our approach relies on assumptions, which are only approximately correct and difficult to test in practice. We assume *causal sufficiency* [22] for the variables in our causal model, which rules out confounding between the variables we intervene on (vaccination times) and the outcome (severe cases). However, we do expect some confounding in practice: at-risk groups like healthcare workers or patients with pre-existing conditions may have a higher incentive to get vaccinated [20]. This could break the assumption of homogeneous subgroups based on age, vaccination status and waning time and lead to overestimated risk factors for the vaccinated  $g(V>0, A)$ .

In our study, we do not model behavioural or policy responses explicitly. A high infection incidence increases perceived risk in the population and prompts voluntary health-protective behaviour [6]. Governments also react to changes in infection incidence through NPIs. Both effects tend to reduce (increase) the effective reproduction number during high (low) incidence periods. Therefore, we expect that our infection and severe-case incidences are overestimated for high-incidence periods and underestimated for low-incidence periods. Our results should be interpreted as counterfactual vaccine allocation scenarios while keeping behavioural and government responses fixed.

Besides approaching the problem of evaluating counterfactual vaccine allocation strategies, this study illustrates a more general problem in causal inference. Causal models typically require joint observations of all relevant variables in order to evaluate counterfactual statements. However, in realistic settings, data is often a limiting factor. In the present study, two crucial parts were not observed: waning times of the severely ill and latent factors related to infection dynamics (*e.g.* base reproduction numbers). By assuming a factorisation for the severity mechanism (2) and incorporating an SEIR-like model together with literature estimates for waning, we were able to address these limitations. In our case, we leveraged domain expertise to merge the different sources of knowledge—namely data, dynamics simulation and literature estimates—into one model. Yet, despite some first steps in this direction [12, 13, 7, 17], merging information from different sources or datasets into a single causal model without strong assumptions stemming from domain knowledge is still an open problem.

Through our simulation-assisted causal model, we show how interactions between different elements of a pandemic such as vaccinations, immunity waning, age-dependent infection spread and risk can be effectively captured. While we have applied our method retrospectively, with parameter inference on observational data, this does not limit its applicability. Each of the sub-modules can be replaced with appropriate assumptions on parts of the system which are yet unknown, as shown in Sec. 3.3. For

example, when COVID-19 vaccines were approved initially, it was not yet clear how long immunity against infection or severe course would last. Thus, besides informing the rollout of COVID-19 vaccination campaigns, we hope that our method can help in future pandemics where the relevant factors may not yet be jointly measured or known from the literature.

## Acknowledgments and Disclosure of Funding

We thank Sebastian Contreras and Simon Bauer for constructive discussions. We thank Franz Paul Spitzner for comments on visualisations. This publication was supported by the German Ministry of Science and Education (BMBF) through the Tübingen AI Center (FKZ 01IS18039A) and by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) through both the Cluster of Excellence “Machine Learning - New Perspectives for Science”, (EXC 2064, project number 390727645) and the project “Cognition of Interaction” (SFB 1528, project number 454648639).

## References

- [1] S. Bauer, S. Contreras, J. Dehning, M. Linden, E. Iftexhar, S. B. Mohr, A. Olivera-Nappa, and V. Priesemann. Relaxing restrictions at the pace of vaccination increases freedom and guards against further COVID-19 waves. *PLOS Computational Biology*, 17(9):e1009288, 2021.
- [2] J. M. Brauner, S. Mindermann, M. Sharma, D. Johnston, J. Salvatier, T. Gavenčiak, A. B. Stephenson, G. Leech, G. Altman, V. Mikulik, A. J. Norman, J. T. Monrad, T. Besiroglu, H. Ge, M. A. Hartwick, Y. W. Teh, L. Chindelevitch, Y. Gal, and J. Kulveit. Inferring the effectiveness of government interventions against COVID-19. *Science*, 371(6531):eabd9338, 2021.
- [3] K. M. Bubar, K. Reinholt, S. M. Kissler, M. Lipsitch, S. Cobey, Y. H. Grad, and D. B. Larremore. Model-informed COVID-19 vaccine prioritization strategies by age and serostatus. *Science*, 371(6532):916–921, 2021.
- [4] N. G. Davies, P. Klepac, Y. Liu, K. Prem, M. Jit, and R. M. Eggo. Age-dependent effects in the transmission and control of COVID-19 epidemics. *Nature Medicine*, 26(8):1205–1211, 2020.
- [5] J. Dehning, J. Zierenberg, F. P. Spitzner, M. Wibral, J. P. Neto, M. Wilczek, and V. Priesemann. Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions. *Science*, 369(6500):eabb9789, 2020.
- [6] P. Dönges, J. Wagner, S. Contreras, E. N. Iftexhar, S. Bauer, S. B. Mohr, J. Dehning, A. Calero Valdez, M. Kretzschmar, M. Mäs, K. Nagel, and V. Priesemann. Interplay between risk perception, behavior, and COVID-19 spread. *Frontiers in Physics*, 10, 2022.
- [7] G. Duarte, N. Finkelstein, D. Knox, J. Mummolo, and I. Shpitser. An automated approach to causal inference in discrete settings. *arXiv preprint arXiv:2109.13471*, 2021.
- [8] R. A. Fisher. *The Design of Experiments*. Oliver and Boyd, 1935.
- [9] S. Flaxman, S. Mishra, A. Gandy, H. J. T. Unwin, T. A. Mellan, H. Coupland, C. Whittaker, H. Zhu, T. Berah, J. W. Eaton, M. Monod, A. C. Ghani, C. A. Donnelly, S. M. Riley, M. A. C. Vollmer, N. M. Ferguson, L. C. Okell, and S. Bhatt. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*, 584(7820):257–261, 2020.
- [10] B. H. Foy, B. Wahl, K. Mehta, A. Shet, G. I. Menon, and C. Britto. Comparing COVID-19 vaccine allocation strategies in India: A mathematical modelling study. *International Journal of Infectious Diseases*, 103:431–438, 2021.
- [11] C. Fraser. Estimating individual and household reproduction numbers in an emerging epidemic. *PLOS ONE*, 2(8):e758, 2007.
- [12] S. H. Garrido Mejia, E. Kirschbaum, and D. Janzing. Obtaining causal information by merging datasets with MAXENT. In *International Conference on Artificial Intelligence and Statistics*, pages 581–603, 2022.
- [13] L. Gresele, J. von Kügelgen, J. Kübler, E. Kirschbaum, B. Schölkopf, and D. Janzing. Causal inference through the structural causal marginal problem. In *Proceedings of the 39th International Conference on Machine Learning*, pages 7793–7824, 2022.

- [14] E. J. Haas, J. M. McLaughlin, F. Khan, F. J. Angulo, E. Anis, M. Lipsitch, S. R. Singer, G. Mircus, N. Brooks, M. Smaja, K. Pan, J. Southern, D. L. Swerdlow, L. Jodar, Y. Levy, and S. Alroy-Preis. Infections, hospitalisations, and deaths averted via a nationwide vaccination campaign using the Pfizer–BioNTech BNT162b2 mRNA COVID-19 vaccine in Israel: a retrospective surveillance study. *The Lancet Infectious Diseases*, 22(3):357–366, 2022.
- [15] S. Han, J. Cai, J. Yang, J. Zhang, Q. Wu, W. Zheng, H. Shi, M. Ajelli, X.-H. Zhou, and H. Yu. Time-varying optimization of COVID-19 vaccine prioritization in the context of limited vaccination capacity. *Nature Communications*, 12(1):4673, 2021.
- [16] P. W. Holland. Statistics and causal inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.
- [17] S. Lee and E. Bareinboim. Causal effect identifiability under partial-observability. In *International Conference on Machine Learning*, pages 5692–5701, 2020.
- [18] L. Matrajt, J. Eaton, T. Leung, and E. R. Brown. Vaccine optimization for COVID-19: Who to vaccinate first? *Science Advances*, 7(6):eabf1374, 2021.
- [19] Ministry of Health Israel. COVID-19 database (licence: Open Definition 2.1). URL <https://data.gov.il/dataset/covid-19>.
- [20] N. Mulberry, P. Tupper, E. Kirwin, C. McCabe, and C. Colijn. Vaccine rollout strategies: The case for vaccinating essential workers early. *PLOS Global Public Health*, 1(10):e0000020, 2021.
- [21] M. Oliu-Barton, B. S. R. Pradelski, Y. Algan, M. G. Baker, A. Binagwaho, G. J. Dore, A. El-Mohandes, A. Fontanet, A. Peichl, V. Priesemann, G. B. Wolff, G. Yamey, and J. V. Lazarus. Elimination versus mitigation of SARS-CoV-2 in the presence of effective vaccines. *The Lancet Global Health*, 10(1):e142–e147, 2022.
- [22] J. Pearl. *Causality*. Cambridge University Press, 2009.
- [23] J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017.
- [24] M. Pollicott, H. Wang, and H. Weiss. Extracting the time-dependent transmission rate from infection data via solution of an inverse ODE problem. *Journal of Biological Dynamics*, 6(2): 509–523, 2012.
- [25] B. Rosen, R. Waitzberg, and A. Israeli. Israel’s rapid rollout of vaccinations for COVID-19. *Israel Journal of Health Policy Research*, 10(1):1–14, 2021.
- [26] D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974.
- [27] U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085, 2017.
- [28] S. Stringhini, A. Wisniak, G. Piumatti, A. S. Azman, S. A. Lauer, H. Baysson, D. De Ridder, D. Petrovic, S. Schrepft, K. Marcus, S. Yerly, I. Arm Vernez, O. Keiser, S. Hurst, K. M. Posfay-Barbe, D. Trono, D. Pittet, L. Gétaz, F. Chappuis, I. Eckerle, N. Vuilleumier, B. Meyer, A. Flahault, L. Kaiser, and I. Guessous. Seroprevalence of anti-SARS-CoV-2 IgG antibodies in Geneva, Switzerland (SEROCoV-POP). *The Lancet*, 396(10247):313–319.
- [29] S. Sunohara, T. Asakura, T. Kimura, S. Ozawa, S. Oshima, D. Yamauchi, and A. Tamakoshi. Effective vaccine allocation strategies, balancing economy with infection control against COVID-19 in Japan. *PLOS ONE*, 16(9):e0257107, 2021.
- [30] S. Y. Tartof, J. M. Slezak, H. Fischer, V. Hong, B. K. Ackerson, O. N. Ranasinghe, T. B. Frankland, O. A. Ogun, J. M. Zamparo, S. Gray, S. R. Valluri, K. Pan, F. J. Angulo, L. Jodar, and J. M. McLaughlin. Effectiveness of mRNA BNT162b2 COVID-19 vaccine up to 6 months in a large integrated health system in the USA: a retrospective cohort study. *The Lancet*, 398(10309):1407–1416, 2021.
- [31] A. R. Tuite, D. N. Fisman, J. C. Kwong, and A. L. Greer. Optimal pandemic influenza vaccine allocation strategies for the Canadian population. *PLOS ONE*, 5(5):e10520, 2010.
- [32] C. Viboud, J. Eisenstein, A. H. Reid, T. A. Janczewski, D. M. Morens, and J. K. Taubenberger. Age- and sex-specific mortality associated with the 1918–1919 influenza pandemic in Kentucky. *The Journal of Infectious Diseases*, 207(5):721–729, 2013.

- [33] T. N. Vilches, S. M. Moghadas, P. Sah, M. C. Fitzpatrick, A. Shoukat, A. Pandey, and A. P. Galvani. Estimating COVID-19 infections, hospitalizations, and deaths following the US vaccination campaigns during the pandemic. *JAMA Network Open*, 5(1):e2142725, 2022.

## A Immunity waning curve

The vaccine efficacy against infection as a function of time since the administration of the second dose of the BionTech vaccine is reported in [30] for discrete time periods up to 6 months. To this data, we fit a logistic curve that tends towards zero efficacy as time increases. For the waning curves after 1 and 3 doses, we assume the same functional relationship as for the second dose, but scale the function such that under full protection the efficacy is 75% and 95%, respectively (efficacy under full protection is around 90% after two doses).

We can use this to derive the waning function  $h^v(W)$ . Let  $VE^v(w)$  be the vaccine efficacy  $w$  weeks after receiving the  $v^{\text{th}}$  dose:

$$VE^v(w) = 1 - \frac{P(I=1|V=v, A=a, T, W=w)}{P(I=1|V=0, A=a, T, W=w)} \quad (12)$$

where we assume that the efficacy against infection is the same for all age groups and constant over time  $T$ . First, we observe that the severe-case probability can be separated into two processes:

$$P_{\tilde{\pi}}(S=1|V, A, T, W) = P_{\tilde{\pi}}(I=1|V, A, T, W) P(S=1|V, A, I=1) \quad (13)$$

(a) the probability of being infected and (b) the probability of having a severe case once infected. We assume that the probability of having a severe case once infected only depends on  $V$  and  $A$  and the immunity against severe courses does not wane over time, as reported in [30]. Then note

$$\frac{P(S=1|V=v, a, t, w)}{P(S=1|V=0, a, t, w)} = \frac{P(S=1|V=v, a, I=1)}{P(S=1|V=0, a, I=1)} \underbrace{\frac{P(I=1|V=v, a, t, w)}{P(I=1|V=0, a, t, w)}}_{1-VE^v(w)} \quad (14)$$

$$= \frac{f(t)g(v, a)h^v(w)}{f(t)g(0, a)} = \frac{g(v, a)}{g(0, a)} h^v(w) \quad (15)$$

$$\Rightarrow h^v(w) = \frac{g(0, a)}{g(v, a)} \frac{P(S=1|V=v, a, I=1)}{P(S=1|V=0, a, I=1)} (1 - VE^v(w)). \quad (16)$$

By definition we have  $h^v(0) = 1$  and combining the above expression for  $W = w \geq 0$  and  $W = 0$  we get

$$\frac{h^v(w)}{h^v(0)} = \frac{1 - VE^v(w)}{1 - VE^v(0)} \Rightarrow h^v(w) = \frac{1 - VE^v(w)}{1 - VE^v(0)}. \quad (17)$$

## B Estimating the infection dynamics correction factor

For the estimation of  $f_{\tilde{\pi}}^1(A, T)$ , we consider the following: let  $P_{\pi}(S=1|V, A, T, W)$  be the severity mechanism under the observed vaccine allocation strategy  $\pi$  and  $P_{\tilde{\pi}}(S=1|V, A, T, W)$  under the post-intervention vaccine allocation strategy  $\tilde{\pi}$ . Then, using eq. (13),

$$\frac{P_{\tilde{\pi}}(S=1|V, A, T, W)}{P_{\pi}(S=1|V, A, T, W)} = \frac{P_{\tilde{\pi}}(I=1|V, A, T, W) P(S=1|V, A, I=1)}{P_{\pi}(I=1|V, A, T, W) P(S=1|V, A, I=1)} = \frac{P_{\tilde{\pi}}(I=1|V, A, T, W)}{P_{\pi}(I=1|V, A, T, W)} \quad (18)$$

where we have used that the process of going from infected to severely ill (??) does not depend on the vaccine allocation strategy. From the factorisation of the severity mechanism (2) and using  $f_{\tilde{\pi}}^1(A, T)=1$  it follows that

$$\frac{P_{\tilde{\pi}}(S=1|V, A, T, W)}{P_{\pi}(S=1|V, A, T, W)} = f_{\tilde{\pi}}^1(A, T). \quad (19)$$

The assumption that  $f^1$  only depends on  $A$  and  $T$  means we assume that the change in infection probability due to the infection dynamics is independent of  $V$  and  $W$ . Hence, the correction factor is the relative change in weekly infection probability for each age group under the counterfactual vaccine allocation strategy:

$$f_{\tilde{\pi}}^1(A, T) = \frac{P_{\tilde{\pi}}(I=1|V, A, T, W)}{P_{\pi}(I=1|V, A, T, W)}. \quad (20)$$

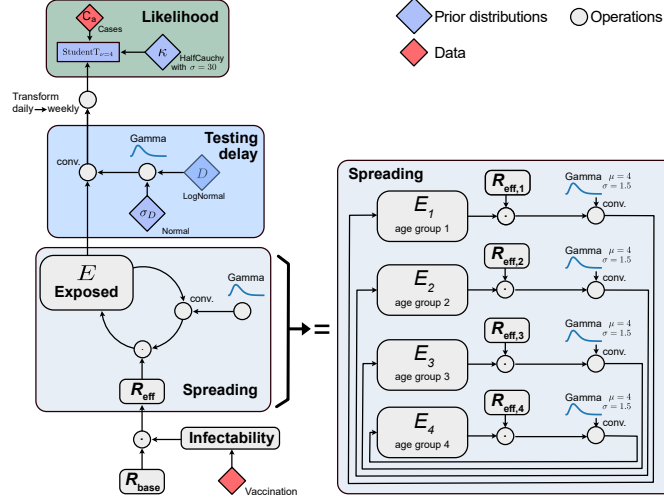


Figure 6: **Overview of the infection dynamics model.** The dynamics model infers the effective  $R_{\text{eff}}$  and base reproduction number as a function of time given the factual vaccine allocation strategy and the number of observed cases  $C_a$ . After inference, the model is used to predict the number of cases under a counterfactual strategy. The right side corresponds to equation (21).

## C Modelling infection dynamics

To estimate the effect of changing the vaccine allocation strategy on the infection dynamics ( $P_{\tilde{\pi}}(I=1|V, A, T, W)$ ) we first infer the parameters of a Bayesian SEIR-like model under the observed strategy  $\pi$ . Then we rerun the model with the inferred reproduction numbers under the counterfactual strategy  $\tilde{\pi}$  to obtain an estimate of the number of infections and the correction factor  $f_{\tilde{\pi}}^1(A, T)$ .

### C.1 The SEIR-like dynamics

In our model each age group  $a$  has its own compartment, and the dynamics follows a discrete renewal process [11, 9, 2] determined by the reproduction number  $R_{\text{eff},a}(t_{\text{day}})$ . These dynamics are discretized with a one-day step, to allow for enough resolution to model the generation interval. The data is only available on a weekly basis, which will require later to sum the cases over one week. We fit one reproduction number per age group. The infections from one age group to another are encoded by a contact matrix  $\mathbf{C}$ . The latent period is modelled by a kernel  $g(\tau)$  which is normalized to 1:

$$E_a(t_{\text{day}}) = \sum_{a'=1}^9 \sqrt{R_{\text{eff},a}(t_{\text{day}})} C_{a,a'} \sqrt{R_{\text{eff},a'}(t_{\text{day}})} \sum_{\tau=0}^{10} E_{a'}(t_{\text{day}} - 1 - \tau) \tilde{g}(\tau) + h_a(t_{\text{day}}), \quad (21)$$

$$S_a(t_{\text{day}}) = S_a(t_{\text{day}} - 1) - E_a(t_{\text{day}} - 1), \quad (22)$$

$$\tilde{g}(\tau) = \text{Gamma}(\tau; \mu=4, \sigma=1.5). \quad (23)$$

$E_a(t_{\text{day}})$  is the number of newly exposed people on day  $t_{\text{day}}$ , who will later become infectious; it is implicitly modelled by the latent kernel  $\tilde{g}$ .  $S_a(t_{\text{day}})$  is the number of susceptible people and  $h_a(t_{\text{day}})$  is an external influx (see below).

We assume that changes in the reproduction number symmetrically affect both the infectiousness and infectability. This is achieved by multiplying the square-root of the reproduction number to the contact matrix from both sides in equation (21).

### C.2 The contact matrix

The contact matrix models the amount of contacts between age groups. There are two extreme cases:

1. There are no contacts between age groups. In this case, the contact matrix would simply be the identity matrix:  $\mathbf{C} = \mathbf{1}$ .

2. The contact between different age groups is the same as within age groups: all-to-all connectivity. Let  $D^{(A=a)}$  be the population of age group  $a$  and  $D = \sum_a D^{(A=a)}$  be the total population. Then, in this case, contact matrix would be  $\mathbf{C} = \vec{\rho} \cdot \vec{1}^T$ , where  $\vec{1}^T = (1, 1, \dots, 1)$  and  $\vec{\rho}^T = (\rho_1, \rho_2, \dots) = \left( \frac{D^{(A=1)}}{D}, \frac{D^{(A=2)}}{D}, \dots \right)$  is the vector of the population share in each age group.

The reality lies somewhere between these two cases. A reasonable interpolation between these extremes should ensure that the largest eigenvector of the contact matrix stays  $\vec{\rho}$  and that the largest eigenvalue is 1. This ensures that under a reproduction number of 1 the case numbers are constant and proportional to the relative population  $\vec{\rho}$  of each age group. These requirements are met by the following matrix:

$$\begin{aligned} \mathbf{C} &= (1 - \gamma) \mathbb{1} + \gamma \vec{\rho} \cdot \vec{1}^T \\ &= \begin{pmatrix} \gamma \rho_1 + (1 - \gamma) & \gamma \rho_1 & \gamma \rho_1 & \dots \\ \gamma \rho_2 & \gamma \rho_2 + (1 - \gamma) & \gamma \rho_2 & \dots \\ \gamma \rho_3 & \gamma \rho_3 & \gamma \rho_3 + (1 - \gamma) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \end{aligned} \quad (24)$$

The parameter  $\gamma$  determines the mixing between age group: at 1 we have all-to-all connectivity, at 0 there is no contact between age groups.

### C.3 The external influx

To account for some infections occurring due to infected travellers entering Israel, we add a random number of infections distributed over each week:

$$h_a(t_{\text{day}}) = h_a^*(t = k_{\text{week}}(t_{\text{day}}))/7, \quad (25)$$

$$h_a^*(t) \sim \text{Weibull}(\lambda = 0.1 \cdot \text{pop}_a/10^6, k = 0.3) \quad \forall a, \forall t, \quad (26)$$

where  $t$  is indexing the weeks included in our analysis,  $k_{\text{week}}(t_{\text{day}})$  is mapping a day  $t_{\text{day}}$  to the corresponding week  $t$ . We chose a Weibull distribution because the long tails allow the occurrence of eventual mass spreading events. The parameters of the Weibull distribution are chosen such that on average 0.1 infections per million inhabitants per day occur from external influx. This is about a fifth of the lowest incidence during the analysis period. The median of the distribution is only at 0.003 infections per million inhabitants per day because of the long tails of the distribution.

### C.4 The effective reproduction number

The effective reproduction number  $R_{\text{eff},a}(t_{\text{day}})$  depends on (1) a base reproduction number  $R_{\text{base},a}(t_{\text{day}})$ , which encodes the amount of social-distancing at time  $t_{\text{day}}$  and (2) on an infectability term  $\text{Infectability}_a(t_{\text{day}})$  which encodes the acquired immunity of the susceptible population of age group  $a$ .

$$R_{\text{eff},a}(t_{\text{day}}) = R_{\text{base},a}(t_{\text{day}}) \cdot \text{Infectability}_a(t_{\text{day}}) \quad (27)$$

### C.5 The infectability

The  $\text{Infectability}_a$  is the fraction of reduced spread because of the acquired immunity in age group  $a$ . It is 1 in a completely non-immune population and reduces with vaccination. Concretely, it is modelled as:

$$\begin{aligned} \text{Infectability}_a(t_{\text{day}}) &= \text{Unv}_a(t) + \text{Vacc1}_a(t) (1 - \mu_{1,a} \mathbf{W1}_{\text{eff},a}(t)) \\ &\quad + \text{Vacc2}_a(t) (1 - \mu_{2,a} \mathbf{W2}_{\text{eff},a}(t)) \\ &\quad + \text{Vacc3}_a(t) (1 - \mu_{3,a} \mathbf{W3}_{\text{eff},a}(t)) \end{aligned} \quad (28)$$

with  $t = k_{\text{week}}(t_{\text{day}})$ , (29)

where  $\text{Unv}_a(t)$ ,  $\text{Vacc1}_a(t)$ ,  $\text{Vacc2}_a(t)$  and  $\text{Vacc3}_a(t)$  are the fractions of unvaccinated, once-, twice- and three-times-vaccinated, respectively, for each age group. Here  $t$  is the week corresponding to the

day  $t_{\text{day}}$ .  $\mu$  denotes the corresponding protection against infection ( $0 < \mu < 1$ ,  $\mu = 0$  corresponds to no protection). These vaccination fractions are modelled in a weekly manner, because the data is only available on a weekly basis.  $W_{\cdot, \text{eff}, a}(t)$  denotes the effective group-wide waning of the immunity, it is 1 if the whole age group had been vaccinated a few days ago and decreases with time.

The waning immunity is calculated by averaging over time the newly vaccinated people and the factor by which their immunity became reduced:

$$W1_{\text{eff}, a}(t) = \frac{\sum_{\tau=0}^t n_{\text{Vacc}1_{a,t}}(\tau) W1(t - \tau)}{\sum_{\tau=0}^t n_{\text{Vacc}1_{a,t}}(\tau)} \quad (30)$$

$$\dots, \quad (31)$$

where  $n_{\text{Vacc}1_{a,t}}(\tau)$  are newly vaccinated people at time  $\tau$  which have not received more than one vaccination by time  $t$ . The waning for the other vaccination stages are modelled similarly.

## C.6 The base reproduction number

The base reproduction number  $R_{\text{base}, a}(t_{\text{day}})$  is assumed to be a slowly changing factor as a function of time. It is modelled as a superposition of logistic change points  $\gamma(t_{\text{day}})$  every 10 days, which are parameterised by the transient length of the change points  $l$ , the date of the change point  $d$  and the effect of the change point  $\Delta\gamma^*$ . The subscripts  $n$  denote the discrete enumeration of the change points:

$$R_{\text{base}, a}(t_{\text{day}}) = R_{0,a} \exp\left(\sum_n \gamma_n(t_{\text{day}})\right) \quad (32)$$

$$R_{0,a} \sim \text{LogNormal}(\mu = 1, \sigma = 1) \quad \forall a \quad (33)$$

$$\gamma_{n,a}(t_{\text{day}}) = \frac{1}{1 + e^{-4/l_{n,a} \cdot (t - d_{n,a})}} \cdot \Delta\gamma_{n,a} \quad (34)$$

$$\Delta\gamma_{n,a} \sim \mathcal{N}(\Delta\gamma_{n-1,a}, \sigma_{\Delta\gamma_a}) \quad \forall n, \forall a \quad \text{with } \Delta\gamma_{0,a} = \log R_{0,a} \quad (35)$$

$$\sigma_{\Delta\gamma_a} \sim \text{HalfCauchy}(0.5) \quad \forall a \quad (36)$$

$$l_{n,a} = \log(1 + \exp(l_{n,a}^\dagger)) \quad (37)$$

$$l_{n,a}^\dagger \sim \mathcal{N}(4, 1) \quad \forall n, \forall a \quad (\text{unit is days}) \quad (38)$$

$$d_{n,a} = 10^{\text{th}} \text{ January } 2021 + 21 \cdot n + \Delta d_{n,a} \quad \text{for } n = 0, \dots, 9 \quad (39)$$

$$\Delta d_{n,a} \sim \mathcal{N}(0, 3.5) \quad \forall n, \forall a \quad (\text{unit is days}). \quad (40)$$

## C.7 The likelihood

Next, we want to define the goodness of fit of our model to the sample data. For that, the number of newly exposed people is delayed by 6 days and summed over one week because the case data is available on a weekly basis. The likelihood of that is modelled by a Student-t distribution, which allows for some outliers because of its heavier tails compared to a normal distribution (green box in Fig. 6). The error of the Student-t distribution is proportional to the square root of the number of cases, which corresponds to the scaling of the errors in a Poisson or Negative Binomial distribution:

$$\hat{C}_a(t) = \sum_{t_{\text{day}}=6}^{13} E(k_{\text{week}}^{-1}(t) - t_{\text{day}}) \quad (41)$$

$$C_a(t) \sim \text{StudentT}_{\nu=4}\left(\mu = \hat{C}_a(t), \sigma = \kappa \sqrt{\hat{C}_a(t) + 1}\right) \quad \text{with} \quad (42)$$

$$\kappa \sim \text{HalfCauchy}(\sigma = 30). \quad (43)$$

Here  $C_a(t)$  is the measured number of weekly cases in the population of age  $a$  as reported by the health authorities, whereas  $\hat{C}_a(t)$  is the modelled number of cases on week  $t$ .  $k_{\text{week}}^{-1}(t)$  transforms the week  $t$  to the first day of the week.



## D Generating vaccine allocation strategies

### D.1 Factual strategy

In the available dataset [19] the vaccination times for first, second and third doses are given for each age group, *i.e.*

$$P(T_i|A) \quad \text{for } i = 1, 2, 3. \quad (44)$$

However, this alone does not uniquely determine the distribution of waning times. To fully specify required joint distributions  $P(T_1, T_2|A)$  and  $P(T_3|T_2, A)$  we employ a greedy algorithm to

1. Maximise the number of vaccine recipients that receive their second dose 3 weeks after their initial dose (or as close to 3 weeks as possible, but not sooner).
2. Have a minimum gap of 12 weeks between second and third dose (constraint).
3. Satisfy (44) (constraint).

### D.2 Uniform strategy

For all counterfactual strategies we require the number of first, second and third doses to match the factual—but not in each age group. In other words,

$$\tilde{P}(T_i) = P(T_i) \quad \text{for } i = 1, 2, 3. \quad (45)$$

where the tilde indicates probabilities in the counterfactual scenario. For the Uniform strategy, we require all age groups to have the same vaccination time distributions:

$$\tilde{P}(T_1, T_2|A) = \tilde{P}(T_1, T_2) \quad \text{and} \quad \tilde{P}(T_3|T_2, A) = \tilde{P}(T_3|T_2). \quad (46)$$

In summary, the greedy algorithm to determine  $\tilde{P}(T_1, T_2|A)$  and  $\tilde{P}(T_3|T_2, A)$  is adapted as follows:

1. Maximise the number of vaccine recipients that receive their second dose 3 weeks after their initial dose (or as close to 3 weeks as possible, but not sooner).
2. Have a minimum gap of 12 weeks between second and third dose (constraint).
3. Satisfy (45) (constraint).
4. Satisfy (46) (constraint).

### D.3 Ranked strategies

In the ranked strategies `ElderlyFirst`, `YoungFirst`, `RiskRanked` and `RiskRankedReversed` all age groups are ranked and vaccinations are assigned to the highest-ranked age group its vaccine uptake rate is reached. The counterfactual vaccine uptake rate per age group is the factual one for the first and second doses, *i.e.*

$$\sum_{t_i=1}^M \tilde{P}(t_i|a) = \sum_{t_i=1}^M P(t_i|a) \quad \text{for } i = 1, 2 \text{ and } \forall a. \quad (47)$$

The vaccine uptake rate for booster shots is slightly relaxed by 2.5% to meet other consistency constraints:

$$\sum_{t_3=1}^M \tilde{P}(t_3|a) = 0.025 + \sum_{t_3=1}^M P(t_3|a) \quad \forall a. \quad (48)$$

The greedy algorithm is adapted to

1. Prioritise age groups according to ranking.
2. Maximise the number of vaccine recipients that receive their second dose 3 weeks after their initial dose (or as close to 3 weeks as possible, but not sooner).
3. Have a minimum gap of 12 weeks between second and third dose (constraint).
4. Satisfy (45) (constraint).
5. Satisfy (47) and (48) (constraint).

## E Target function

Our goal is to compute the expected number of severe cases after intervening on the distribution of vaccination times:

$$s(\tilde{\pi} = \tilde{P}(T_1, T_2, T_3|A)) := MD \mathbb{E}[S|\text{do}(T_1, T_2, T_3 \sim \tilde{P}(T_1, T_2, T_3|A))] \quad (49)$$

where  $\tilde{P}(T_1, T_2, T_3|A)$  is the distribution of vaccination times in the counterfactual scenario. Using results for stochastic policies from [22], repeated application of the law of total probabilities (TP) and exploiting conditional independences (CI) implied by the causal graph (see Fig. 1) we can derive an expression for (49) in terms of known conditional probabilities:

$$\begin{aligned} & s(\tilde{\pi} = \tilde{P}(T_1, T_2, T_3|A)) \\ & := MD \mathbb{E}[S|\text{do}(T_1, T_2, T_3 \sim \tilde{P}(T_1, T_2, T_3|A))] \end{aligned} \quad (50)$$

$$= MD P\left(S=1|\text{do}(T_1, T_2, T_3 \sim \tilde{P}(T_1, T_2, T_3|A))\right) \quad (51)$$

$$\stackrel{[22, \text{Ch. 4.2}]}{=} MD \sum_a P(a) \sum_{t_1, t_2, t_3=1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) P(S=1|\text{do}(t_1, t_2, t_3), a) \quad (52)$$

$$\begin{aligned} & \stackrel{\text{TP}}{=} MD \sum_a P(a) \sum_{t_1, t_2, t_3=1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) \sum_{t=1}^M \underbrace{P(t|\text{do}(t_1, t_2, t_3), a)}_{\stackrel{\text{CI}}{=} P(t)} \\ & P(S=1|\text{do}(t_1, t_2, t_3), a, t) \end{aligned} \quad (53)$$

$$\begin{aligned} & \stackrel{\text{TP}}{=} MD \sum_a P(a) \sum_{t=1}^M \underbrace{P(t)}_{=1/M} \sum_{t_1, t_2, t_3=1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) \\ & \sum_{w=1}^M \underbrace{P(w|\text{do}(t_1, t_2, t_3), a, t)}_{\stackrel{\text{CI}}{=} P(w|\text{do}(t_1, t_2, t_3), t)} P(S=1|\text{do}(t_1, t_2, t_3), a, t, w) \end{aligned} \quad (54)$$

$$\begin{aligned} & \stackrel{\text{TP}}{=} D \sum_a P(a) \sum_{t=1}^M \sum_{t_1, t_2, t_3=1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) \sum_{w=1}^M P(w|\text{do}(t_1, t_2, t_3), t) \\ & \sum_{v=0}^3 \underbrace{P(v|\text{do}(t_1, t_2, t_3), a, t, w)}_{\stackrel{\text{CI}}{=} P(v|\text{do}(t_1, t_2, t_3), t)} \underbrace{P(S=1|\text{do}(t_1, t_2, t_3), v, a, t, w)}_{\stackrel{\text{CI}}{=} P(S=1|v, a, t, w)} \end{aligned} \quad (55)$$

$$\begin{aligned} & = D \sum_a P(a) \sum_{t=1}^M \sum_{t_1, t_2, t_3=1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) \\ & \sum_{w=1}^M P(w|\text{do}(t_1, t_2, t_3), t) \sum_{v=0}^3 P(v|\text{do}(t_1, t_2, t_3), t) P(S=1|v, a, t, w) \end{aligned} \quad (56)$$

$$\begin{aligned} & = D \sum_a P(a) \sum_{t=1}^M \sum_{t_1, t_2, t_3=1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) \\ & \sum_{w=1}^M P(w|t_1, t_2, t_3, t) \sum_{v=0}^3 P(v|t_1, t_2, t_3, t) P(S=1|v, a, t, w). \end{aligned} \quad (57)$$

In the last step we use that since  $\{T_1, T_2, T_3, T\}$  are the parents of  $V$  and  $W$  we can replace the do-operators by conditional probabilities [22]. Also note that for notational convenience, we set the vaccination time  $t_i = M + 1$  for a patient who has not received the  $i^{\text{th}}$  dose during the considered time window  $t \in \{1, \dots, M\}$ .

The waning time  $w$  (number of weeks since the last dose was received) depends deterministically on the vaccination times  $t_1, t_2, t_3$ :

$$P(w|t_1, t_2, t_3, t) = \begin{cases} 1 & \text{if } w = \max_{i \in \{1,2,3\}} [t - t_i]^+, \\ 0 & \text{else,} \end{cases} \quad \text{where } [x]^+ = \begin{cases} x & \text{if } x \geq 0, \\ 0 & \text{else.} \end{cases} \quad (58)$$

The vaccination status  $v$  also depends deterministically on the vaccination times  $t_1, t_2, t_3$ :

$$P(v|t_1, t_2, t_3, t) = \begin{cases} 1 & \text{if } v = \max_{i \in \{0,1,2,3\}} (i \cdot \text{sgn}(t - t_i + 1)), \\ 0 & \text{else,} \end{cases} \quad (59)$$

where  $\text{sgn}$  is the sign function. These deterministic relationships can be used to eliminate the corresponding conditionals from the target function:

$$\begin{aligned} & s(\tilde{\pi} = \tilde{P}(T_1, T_2, T_3|A)) \\ &= D \sum_a P(a) \sum_{t=1}^M \\ & \quad \times \left[ \sum_{t_1=t+1}^{M+1} \sum_{t_2=t+1}^{M+1} \sum_{t_3=t+1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) P(S=1|t, a, v=0, w=0) \right. \\ & \quad + \sum_{t_1=1}^t \sum_{t_2=t+1}^{M+1} \sum_{t_3=t+1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) P(S=1|t, a, v=1, w=(t-t_1)) \\ & \quad + \sum_{t_1=1}^t \sum_{t_2=1}^t \sum_{t_3=t+1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) P(S=1|t, a, v=2, w=(t-t_2)) \\ & \quad \left. + \sum_{t_1=1}^t \sum_{t_2=1}^t \sum_{t_3=1}^t \tilde{P}(t_1, t_2, t_3|a) P(S=1|t, a, v=3, w=(t-t_3)) \right]. \quad (60) \end{aligned}$$

Substituting the factorisation for the severity mechanism  $P(S=1|v, a, t, w)$  we get:

$$\begin{aligned} & s(\tilde{\pi} = \tilde{P}(T_1, T_2, T_3|A)) \\ &= D \sum_a P(a) \sum_{t=1}^M f^0(t) f_{\tilde{\pi}}^1(a, t) \\ & \quad \times \left[ \sum_{t_1=t+1}^{M+1} \sum_{t_2=t+1}^{M+1} \sum_{t_3=t+1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) g(0, a) \right. \\ & \quad + \sum_{t_1=1}^t \sum_{t_2=t+1}^{M+1} \sum_{t_3=t+1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) g(1, a) h^1(t-t_1) \\ & \quad + \sum_{t_1=1}^t \sum_{t_2=1}^t \sum_{t_3=t+1}^{M+1} \tilde{P}(t_1, t_2, t_3|a) g(2, a) h^2(t-t_2) \\ & \quad \left. + \sum_{t_1=1}^t \sum_{t_2=1}^t \sum_{t_3=1}^t \tilde{P}(t_1, t_2, t_3|a) g(3, a) h^3(t-t_3) \right]. \quad (61) \end{aligned}$$

## F Severe-case incidence over time

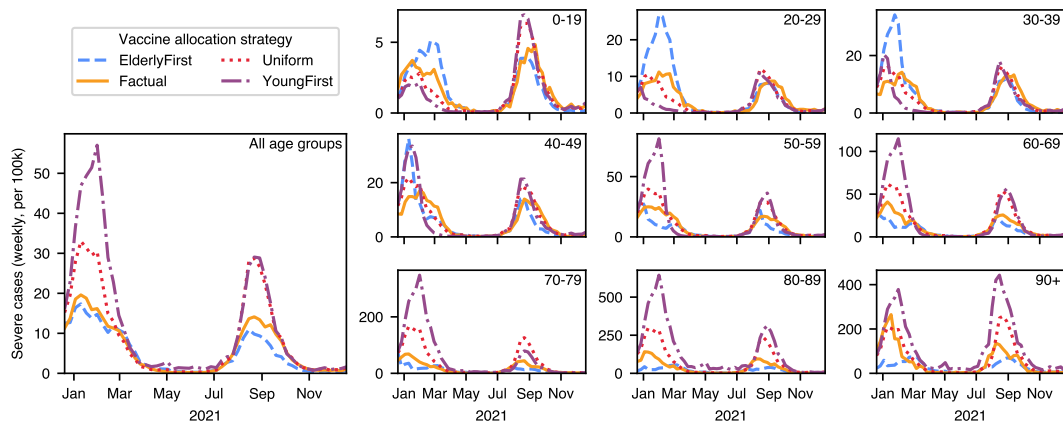


Figure 7: Expected weekly number of severe-case incidence for the entire population (left) and in each age group (right) for the factual and counterfactual vaccine allocation strategies. The right panels show the trade-offs in severe-case incidence between age groups under different vaccine allocation strategies.