MODELING BETWEEN-POPULATION VARIATION IN COVID-19 DYNAMICS IN HUBEI, LOMBARDY, AND NEW YORK CITY

Bryan Wilder[1], Marie Charpignon[2], Jackson A. Killian[1], Han-Ching Ou[1], Aditya Mate[1], Shahin Jabbari[1], Andrew Perrault[1], Angel Desai[3], Milind Tambe[1], Maimuna S. Majumder[4,5]

[1]School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA
[2]MIT Institute for Data, Systems, and Society, Cambridge, MA, USA
[3]International Society for Infectious Diseases, Brookline, MA, USA
[4]Department of Pediatrics, Harvard Medical School, Boston, MA, USA
[5]Computational Health Informatics Program, Boston Children's Hospital, Boston, MA, USA

May 28, 2020

ABSTRACT

As the COVID-19 pandemic continues, formulating targeted policy interventions that are informed by differential SARS-CoV2 transmission dynamics will be of vital importance to national and regional governments. We develop an individual-level model for SARS-CoV2 transmission that accounts for location-dependent distributions of age, household structure, and comorbidities. We use these distributions together with age-stratified contact matrices to instantiate specific models for Hubei, China; Lombardy, Italy; and New York City, United States. Using data on reported deaths to obtain a posterior distribution over unknown parameters, we infer differences in the progression of the epidemic in the three locations. We also examine the role of transmission due to particular age groups on total infections and deaths. The effect of limiting contacts by a particular age group varies by location, indicating that strategies to reduce transmission should be tailored based on population-specific demography and social structure. These findings highlight the role of between-population variation in formulating policy interventions. Across the three populations though, we find that targeted "salutary sheltering" by 50% of a single age group may substantially curtail transmission when combined with the adoption of physical distancing measures by the rest of the population.

Introduction

Since December 2019, the COVID-19 pandemic – propagated by the novel coronavirus, SARS-CoV2 – has resulted in significant morbidity and mortality [1]. As of May 19, 2020, an estimated 4,800,000 individuals have been infected, with over 318,000 fatalities worldwide [2]. Key factors such as existing comorbidities and age appear to play a role in an increased risk of mortality [3]. Epidemiological studies have provided significant insights into the disease and its transmission dynamics to date [4–7]. However, as national and regional governments begin to implement broad-reaching policies in response to rising case counts and stressed healthcare systems, tailoring these polices based on an understanding of how population-specific demography impacts outbreak dynamics will be vital. Previous modeling studies have not incorporated the rich set of household demographic features needed to address such questions.

This study develops a stochastic agent-based model for SARS-CoV2 transmission which accounts for distributions of age, household types, comorbidities, and contact between different age groups in a given population (Fig. 1). Our model accounts for both within-household contact (simulated via household distributions taken from census data) and out-of-household contact using age-stratified, country-specific estimated contact matrices [8]. We instantiate the model for Hubei, China; Lombardy, Italy; and New York City, United States, developing a Bayesian inference strategy for estimating the distribution of unknown parameters using data on reported deaths in each location. This enables us to uncover differences in the progression of the epidemic in each location. We also examine how transmission by particular age groups contributes to infections and deaths in each location, allowing us to compare the efficacy of efforts to reduce transmission across said groups. There is large between-population variation in the role played by any individual age group. However, across populations, both infections and deaths are substantially reduced by a combination of population-wide physical distancing and "salutary sheltering" – a term we coin here to describe

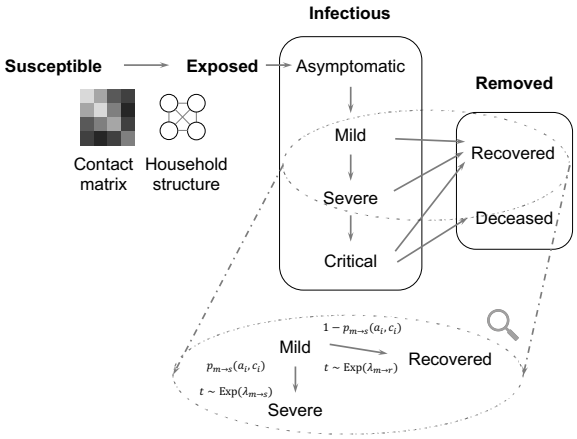

Figure 1: We use a modified SEIR model, where the infectious states are subdivided into levels of disease severity. The transitions are probabilistic and there is a time lag for transitioning between states. For example, the magnified section shows the details of transitions between mild, recovered, and severe states. Each arrow consists of the probability of transition (e.g., $p_{m\to s}(a_i, c_i)$ denotes the probability of progressing from mild to severe) as well as the associated time lag (e.g., the time t for progression from mild to severe is drawn from an exponential distribution with mean $\lambda_{m\to s}$). a_i and c_i denote the age and set of comorbidities for the infected individual i.

individuals who shelter in place irrespective of their exposure or infectious state – by half the individuals in a specific age group, without the need for potentially untenable policies such as indefinite sheltering of all older adults.

Results

Inferring differences in dynamics between populations

Using our model, we estimate posterior distributions over unobserved quantities which characterize the dynamics of the epidemic in a particular location. This section presents estimates for two quantities: first, the basic reproduction number r_0, and second, the rate at which infections are documented. Neither quantity is directly observable in the data due to substantial underdocumentation of infections; however, these estimates are needed to characterize the scope of the outbreak in a particular location, the degree to which existing testing strategies capture new infections, and the rate at which infections are expected to increase in the absence of any intervention. These findings are critical to formulate policy interventions that are tailored to the outbreak as it evolves in a given population. We start by providing a brief overview of our inference strategy and model validation and then present the main estimates.

There are three model parameters for which values are not precisely estimated in the literature. Each such parameter is instead drawn from a prior distribution. First is p_{inf}, the probability of infection given contact with an infected individual. This determines the level of transmissibility of the disease. Second is t_0, the start time of the infection, which is not precisely characterized in most locations and has an impact due to rapid doubling times. Third is a parameter d_{mult}, which accounts for differences in mortality rates between locations that are *not* captured by demographic factors in the model (e.g., the impact of variation in health system capacities). d_{mult} is a multiplier to the baseline mortality rate from [9] and is applied uniformly across age groups. We also include an age-specific multiplier to the mortality rate for individuals over 60 in Lombardy, which is calibrated independently of the other parameters to match the fraction of deaths attributed to the 60+ age group (which is significantly higher in Lombardy than the other two locations [9–11]). Details of the prior distributions and the modeled scenario in each location can be found in the appendix. We incorporate reduced person-to-person contact after mobility restrictions were imposed in each location, basing the strength of the effect on post-lockdown contact surveys [12] and mobility data collected from mobile phones [13].

By conditioning on the observed time series of deaths, we obtain a joint posterior distribution over both the unobserved model states, such as the number of people infected at each time step, as well as the three unknown parameters. We use reported deaths because they are believed to be better documented than infections and perform a sensitivity analysis to account for possible underdocumentation of deaths [14, 15]. Fig. 2 shows that the model closely reproduces the observed time series of deaths in each location. In the appendix (Figs. S1-S3), we also perform out-of-sample validation by fitting the model using a portion of the time series and assessing the accuracy of the predictive posterior distribution on data that was not used to fit the model.

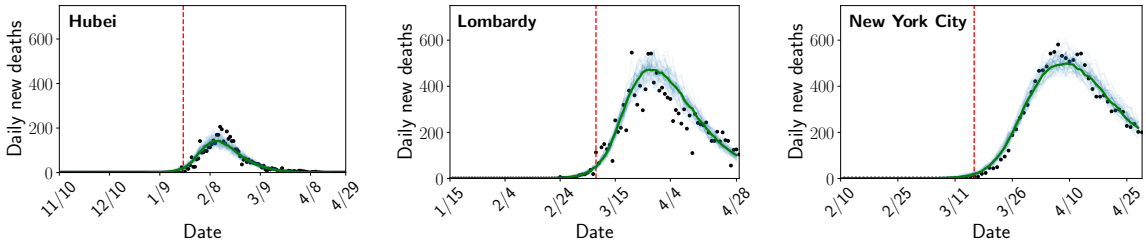

Figure 2: Posterior distribution over the number of deaths each day compared to the number of reported deaths. Light blue lines are individual samples from the posterior, green is the median, and the black dots are the number of reported deaths. The red dashed line represents the start of modeled contact reductions in each location.

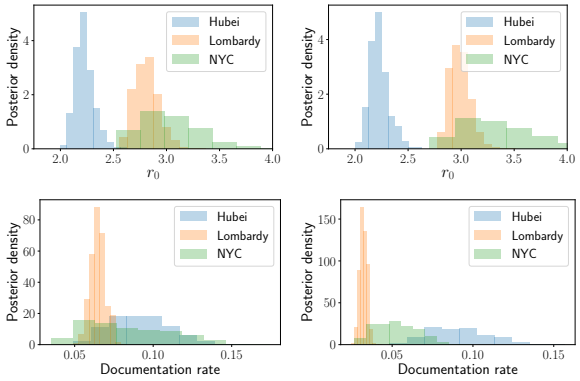

Figure 3: Posterior distribution over r_0 and the fraction of infections documented in each location. Left: conditioning on reported deaths. Right: conditioning on deaths being twice what was reported.

The top row of Fig. 3 shows the posterior distribution over r_0 in each location. Substantial differences are evident between the three locations. The posterior median is 2.21 in Hubei (90% credible interval: 2.10–2.41), 2.80 in Lombardy (2.66–3.01), and 3.06 in New York City (2.65–3.59). The estimates for Hubei fall within the range of a number of existing estimates [16], while the interval for Lombardy is lower than (but overlapping with) the interval 2.9–3.2 estimated by previous work [17]. The estimated r_0 for New York City is larger than either Hubei or Lombardy. To our knowledge, this estimate constitutes the first r_0 assessment in the literature specifically for New York City. These between-population differences are not impacted by a sensitivity analysis for underreporting of deaths, shown in Fig. 3. Death totals from Hubei have been substantially revised upwards to correct for underreporting in the early stages of the epidemic [18], but such corrections are either unavailable or rapidly evolving for Lombardy and New York City. Our sensitivity analysis assumes that deaths in Lombardy and New York City are twice what was reported, consistent with preliminary investigations of excess mortality data [14, 15].

The bottom row of Fig. 3 shows the posterior distribution over the fraction of infections that were documented in each location (obtained by dividing the number of confirmed cases in each location by the number of infections in the simulation under each sample from the posterior). Documentation rates are uniformly low, indicating undocumented infections in all locations; however, we estimate lower documentation in Lombardy (90% credible interval: 5.7–7.3%) than in either New York City (4.8–13.1%) or Hubei (6.5–12.2%).

Although we estimate a substantial number of undocumented infections, all locations remain potentially vulnerable to second-wave outbreaks, with the median percentage of the population infected at 7.5% in Hubei, 11.7% in Lombardy, and 25.2% in New York City. Recent serological surveys have estimated 25% infected in New York City [19], consistent with our distribution. When assuming that deaths are underreported by a factor of two in Lombardy and New York City, the median percentage infected is 23.4% in Lombardy and 38.6% in New York City[1]. Overall, our estimates for r_0 and the remaining population of susceptible individuals indicate that Hubei, Lombardy, and New York City could experience new outbreaks in the absence of continued interventions to reduce transmission. Despite this, between-population differences remain substantial; Hubei, Lombardy, and

[1]Of note, even in a scenario with substantially more deaths than documented, it is possible for the fraction infected to be lower than these estimates. Our model's contact patterns capture the general population, but there is the potential for excess deaths to occur disproportionately in high-risk settings with anomalous contact patterns (e.g., reports have linked a large number of deaths to elder care facilities [20]). In such circumstances, larger death totals would not necessarily indicate a substantial increase in the fraction of the entire population infected.

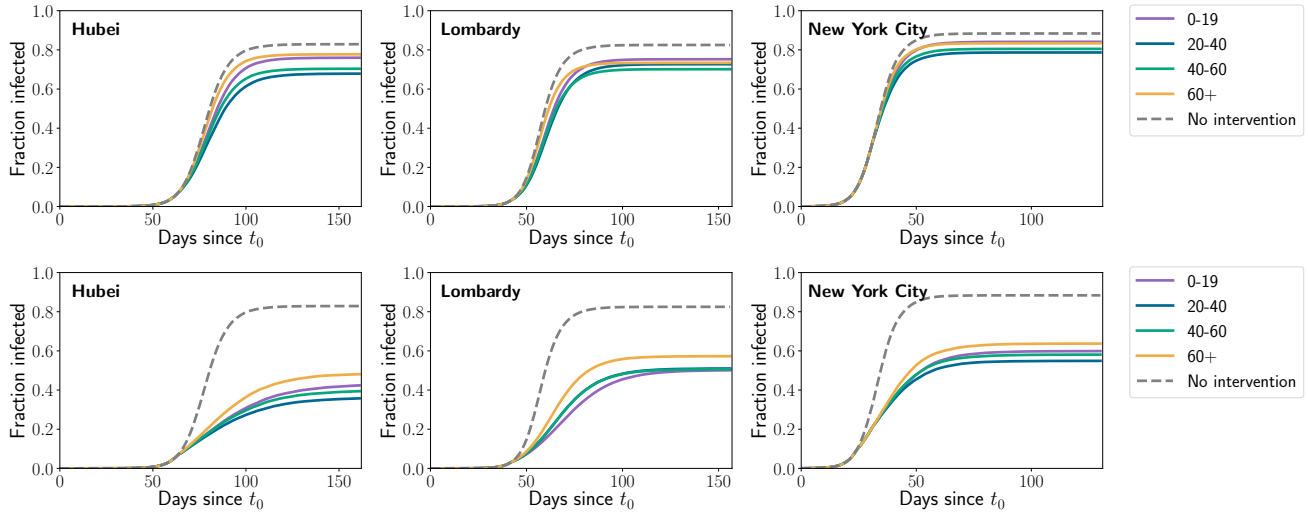

Figure 4: Fraction of the population infected when 50% of each age group shelters at home. Top row: no physical distancing by the rest of the population. Bottom row: physical distancing by the rest of the population. Each line denotes a scenario where the 50% of the indicated age group shelters at home and the rest of the population continues either normal contact patterns (top row) or normal contact subject to physical distancing (bottom row). The "no intervention" line denotes a baseline scenario with no sheltering at home and no physical distancing. Transmission by younger groups generally contributes more towards total infections, but the extent of the differences varies by population.

New York City each have thus far had distinct experiences with COVID-19 that must be considered with respect to future policy responses.

Containment Policies: Salutary Sheltering and Physical Distancing

Various interventions – from complete lockdown to physical distancing recommendations – have been implemented worldwide in response to COVID-19. Within these are a range of alternatives. For example, a government could encourage some percentage of a given age group to remain sheltered in place, while the rest of the population could continue in-person work and social activities. Age-specific policies are particularly relevant because they have already been employed in some countries (e.g., US CDC recommendations that people above 65 years old shelter in place [21]) and because older age groups are more likely to be able to telecommute (at least in the US [22, 23]).

Here, we investigate to what extent the epidemic in each of our three locations of interest can be mitigated by encouraging a single age group to engage in salutary sheltering or whether the entire population must also be asked to adopt physical distancing. Accordingly, we compare two scenarios. First, we simulate salutary sheltering for a fraction of a single age group while leaving the rest of the population's behavior unaltered. Second, we simulate salutary sheltering for a fraction of a single age cohort and physical distancing measures among the rest of the population. We also present a baseline scenario with neither intervention. We model physical distancing as reducing daily contacts by a factor of two for all individuals (who are not sheltering). The supplementary text discusses concrete suggestions to implement such physical distancing in various settings along with details of the experiments described. While this case study applies to Hubei, Lombardy, and New York City, it could be extended to other locations using population-specific demographic data as well.

Fig. 4 shows the fraction of the population infected in each scenario, while Fig. 5 shows the number of resulting deaths. All values are expectations over the posterior distribution of simulation parameters for a given location. Three main results emerge, each of which has implications for the design of targeted non-pharmaceutical interventions.

First, sheltering in place by 50% of a single age group is never sufficient to contain the epidemic by itself. Across all settings, and all choices of which age groups to shelter, at least 60% of the population is infected without physical distancing by the rest of the population, yielding a reduction of at most 20% from the baseline (no intervention). When physical distancing is added however, the percentage infected can be driven below 50%, and sometimes as low as 30%, depending on location and choice of age group to shelter, with possible reductions of 40% or more from baseline. Deaths also decrease commensurately. This highlights that improved protection of a single age group (e.g., the 60+ group) from infection is not independently sufficient to avert large numbers of infections and deaths without physical distancing by the rest of the population.

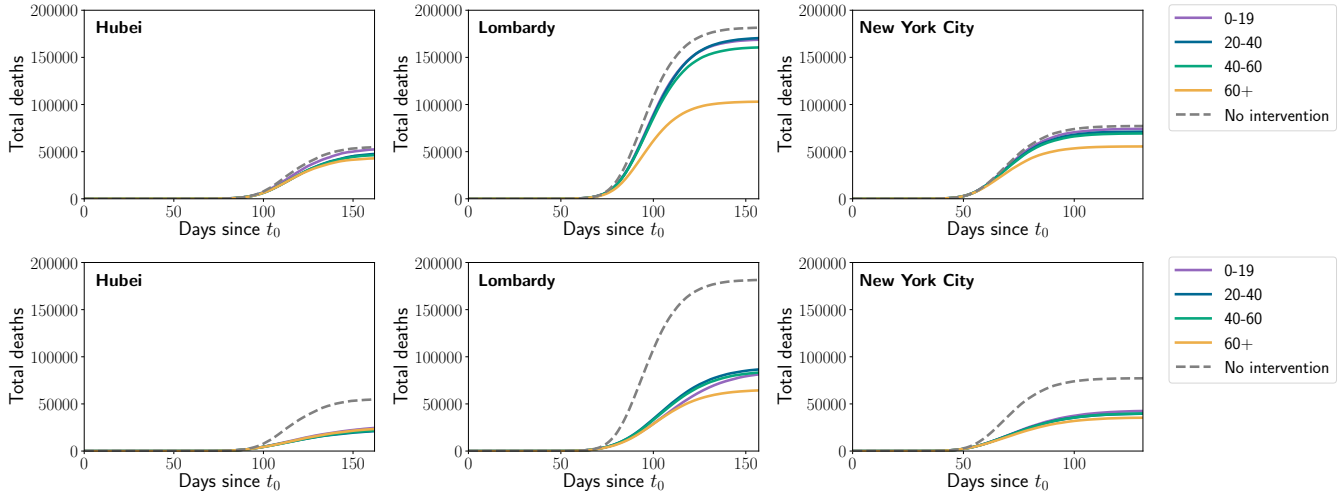

Figure 5: Total number of deaths when 50% of each age group shelters at home. Top row: no physical distancing by the rest of the population. Bottom row: physical distancing by the rest of the population. Each line denotes a scenario where the 50% of the indicated age group shelters at home and the rest of the population continues either normal contact patterns (top row) or normal contact subject to physical distancing (bottom row). The "no intervention" line denotes a baseline scenario with no sheltering at home and no physical distancing. Sheltering by older groups is more effective at reducing the number of deaths without physical distancing, but the adoption of physical distancing by the rest of the population makes sheltering by younger groups as effective in Hubei and nearly as effective in New York City.

Second, the relative effectiveness of sheltering different age groups depends on whether the goal is to minimize infections or deaths. On the whole, younger age groups contribute more towards infections because they tend to have a larger number of daily contacts [8]; however, due to assortative mixing within age groups, these infections tend to be concentrated in other younger people who are at lower risk of fatality. By contrast, older age groups have fewer contacts, meaning that sheltering by these groups has less of an impact on infections; however, since older individuals are at higher risk of fatality, even a small reduction in infections in this group translates into a disproportionate number of deaths averted.

Third, the impact of sheltering by different age groups changes, sometimes substantially, depending on the location under consideration. For example, sheltering by the 60+ age group reduces deaths much more substantially in Lombardy than in either Hubei or New York City because Italian fatalities are concentrated much more heavily in older groups, with 95% of reported deaths in the 60+ age group compared to 80% in Hubei and 74% in New York City [9–11]. In other contexts, the benefits of sheltering the 60+ age group can even be dominated by that of sheltering younger age groups. For example, in the scenario where Hubei adopts population-wide physical distancing, sheltering by the 40-60 age group is more effective at reducing infections than sheltering the 60+ age group while producing a comparable, or slightly lower, number of deaths. Between-population differences in patterns of contact and age-dependent vulnerability to the disease can together alter the effectiveness of sheltering a particular age group.

Building on Hubei, Lombardy, and New York City, our model suggests that hybrid policies that combine targeted salutary sheltering by one sub-population and physical distancing by the rest could be as effective at limiting the final outbreak size as salutary sheltering of an entire sub-population, while preserving social ties and avoiding complete disruption of the economy. Nevertheless, the impact of potential interventions can vary substantially between populations, warranting careful consideration of the demography and social structure of a particular place prior to intervention implementation. Our analysis can be readily extended to other locations by parameterizing our model for a new population using existing demographic data and age-stratified contact patterns, allowing analysis of population-specific intervention.

Discussion and Future Work

In this study, we developed a model of SARS-CoV2 transmission that incorporates household structure, age distributions, comorbidities, and age-stratified contact patterns in Hubei, Lombardy, and New York City and created simulations using available demographic information from these three locations. Our findings suggest that in some locations, substantial reductions in SARS-CoV2 spread can be achieved by less drastic options short of population-wide sheltering in place. Instead, targeted salutary sheltering of specific age groups combined with adherence to physical distancing by the rest of the population may be sufficient to thwart a substantial fraction of infections and deaths. Physical distancing could be achieved by engaging in activities

such as staggered work schedules, increasing spacing in restaurants, and prescribing times to use the gym or grocery store. Specific mechanisms and considerations for implementing physical distancing are documented in the supplementary text. It is important to note that between-population differences in the impact of sheltering different age groups can be substantial. Contact patterns, household structures, and variation in fatality rates (whether due to demographics or factors such as health system capacity) all influence the number of infections or deaths averted by sheltering a particular group. Thus, the implementation of physical distancing and sheltering policies should be tailored to the dynamics of COVID-19 in a particular population.

From a pragmatic perspective, targeted salutary sheltering may not be realistic for all populations. Its feasibility relies on access to safe shelter, which does not reflect reality for all individuals. In addition, sociopolitical realities may render this recommendation more feasible in some populations than in others. Concerns for personal liberty, discrimination against sub-segments of the population, and societal acceptability may prevent the adoption of targeted salutary sheltering in some regions of the world. Allowing salutary sheltering to operate on a voluntary basis using a shift system (rather than for indefinite time periods) may address some of these issues. Future work should formulate targeted recommendations about salutary sheltering and physical distancing by age group or other stratification adapted to a specific country's workforce.

One strength of this study is our ability to assess targeted interventions such as salutary sheltering in a population-specific manner. Existing modeling work of COVID-19 has largely focused on simpler compartmental or branching process models which do not allow for such assessments. While these models have played an important role in estimating key parameters such as r_0 [5, 7] and the rate at which infections are documented [24], as well as in the evaluation of prospective non-pharmaceutical interventions [25, 26], they do not characterize how differences in demography impact the course of an epidemic in a particular location. Our focus on population-specific demography allows for further refinement of current mortality estimates and is a strength of this study. r_0 estimates in this study are generally comparable to other estimates in the literature [16], although our model yields higher estimates for New York City and Lombardy than Hubei – possibly due to differential mask-wearing practices [27] or adoption of behavioral interventions such as hand hygiene [28]. Reporting rates estimated in this study were generally lower than those in prior studies [29], although the trend across locations is consistent. One potential explanation is that Russell et al. estimate documentation from death data using a case fatality rate (CFR) from the literature while our model uses an *infection* fatality rate (IFR). The IFR is lower because it includes all infections, not only those that become confirmed cases. A lower fatality rate in turn implies that each additional infection is less likely to result in death, and so a greater number of total infections are required to account for the observed number of deaths.

One key advantage of our framework is its flexibility. Our model is modifiable to test different policies or simulate additional features with greater fidelity across a variety of populations. Examples of future work that can be accommodated include analysis of contact tracing and testing policies, health system capacity, and multiple waves of infection after lifting physical distancing restrictions. Our model includes the necessary features to simulate these scenarios while remaining otherwise parsimonious, a desirable feature given uncertainties in data reporting.

This study is not without limitations however. While several comorbidities associated with mortality in COVID-19 were accounted for, the availability of existing data limited the incorporation of all relevant comorbidities. Most notably, chronic pulmonary disease was not included although it has been associated with mortality in COVID-19 [30], nor was smoking, despite its prevalence in both China and Italy [31, 32]. Gender-mediated differences were also excluded, which may be important for both behavioral reasons (e.g., adoption of hand-washing [33, 34]) and biological reasons (e.g, the potential protective role of estrogen in SARS-CoV infections [35]). Nevertheless, these factors can all be incorporated into the model as additional data becomes available.

It is also worth noting that we have not yet attempted to model super-spreader events in our existing framework. Such events may have been consequential in South Korea [36], and future work could attempt to model the epidemic there by incorporating a dispersion parameter into the contact distribution, a method which has been employed in other models [5].

Despite these limitations, this study demonstrates the importance of considering population and household demographics when attempting to better define outbreak dynamics for COVID-19. Furthermore, this model highlights potential policy implications for non-pharmaceutical interventions that account for population-specific demographic features and may provide alternative strategies for national and regional governments moving forward.

Materials and Methods

This section provides an overview of our modeling and inference strategy. Additional details may be found in the appendix.

Model

We develop an agent-based model for COVID-19 spread which accounts for the distributions of age, household types, comorbidities, and contact between different age groups in a given population. The model follows a *susceptible-exposed-infectious-removed*

(SEIR) template [37, 38]. Specifically, we simulate a population of n agents (or individuals), each with an age a_i, a set of comorbidities c_i, and a household (a set of other agents). We stratify age into ten-year intervals and incorporate hypertension and diabetes as comorbidities due to their worldwide prevalence [39] and association with higher risk of in-hospital death for COVID-19 patients [3]. However, our model can be expanded to include other comorbidities of interest in the future. The specific procedure we use to sample agents from the joint distribution of age, household structures, and comorbidities may be found in the SI.

The disease is transmitted over a contact structure, which is divided into in-household and out-of-household groups. Each agent has a household consisting of a set of other agents (see the appendix for details on how households are generated using country-specific census information). Individuals infect members of their households at a higher rate than out-of-household agents. We model out-of-household transmission using country-specific estimated contact matrices [8]. These matrices state the mean number of daily contacts an individual of a particular age stratum has with individuals from each of the other age strata. We assume demographics and contact patterns in each location are well-approximated by country-level data.

The model iterates over a series of discrete time steps, each representing a single day, from a starting time t_0 to an end time T. There are two main components to each time step: disease progression and new infections. The progression component is modeled by drawing two random variables for each individual each time they change severity levels (e.g., on entering the mild state). The first random variable is Bernoulli and indicates whether the individual will recover or progress to the next severity level. The second variable represents the amount of time until progression to the next severity level. We use exponential distributions for almost all time-to-event distributions, a common choice in the absence of specific distributional information [40, 41]. The exception is the incubation time between asymptomatic and mild states, where more specific information is available; here, we use a log-normal distribution (see $\mu_{e\to m}$ and $\sigma^2_{e\to m}$ in Table S1) based on estimates by [42]. Table S1 summarizes all distributions and their parameters, and the appendix describes how we estimate age- and comorbidity-dependent severity progression.

In the new infections component, infected individuals infect each of their household members with probability p_h at each time step. p_h is calibrated so that the total probability of infecting a household member before either isolation or recovery matches the estimated secondary attack rate for household members of COVID-19 patients (i.e., the average fraction of household members infected) [43]. Infected individuals draw outside-of-household contacts from the general population using the country-specific contact matrix. For an infected individual of age group i, we sample $w^s_{ij} \sim \text{Poisson}(M^s_{ij})$ contacts for each age group j and setting s where M^s is the country-specific contact matrix for setting s. We include contacts in work, school, and community settings. Poisson distributions are a standard choice for modeling contact distributions [8]. Then, we sample w^s_{ij} contacts of age j uniformly with replacement, and each contact is infected with the probability p_{inf}, the probability of infection given contact. There is evidence to suggest that the probability of infection is higher for an older individual than younger given the same exposure [12], consistent with decline in immune function with age. We adjust for this by letting the probability of infection be βp_{inf} when the exposed individual is over the age of 60, for $\beta > 1$. β is calibrated to match the fraction of deaths in China attributed to individuals over the age of 60, resulting in a value of 1.25. This is consistent with the relationship between age and attack rate amongst close contacts of a confirmed case reported by [12], where the increase in risk of infection for a contact over 65 years old was estimated in the range 1.12–1.92.

Inference of posterior distributions

We infer unknown model parameters and states in a Bayesian framework. This entails placing a prior distribution over the unknown parameters, and then specifying a likelihood function for the observable data, the time series of deaths reported in a location. We posit the following generative model for the observed deaths:

$$
\begin{aligned}
&p_{\text{inf}}, d_{\text{mult}}, t_0 \sim \mathcal{U} \\
&d_1 .... d_T \sim ABM(p_{\text{inf}}, d_{\text{mult}}, t_0) \\
&o_t \sim \text{NegativeBinomial}(d_t, \sigma^2_{obs}) \; t = 1...T
\end{aligned}
$$

where $\mathcal{U}$ denotes a joint uniform prior, ABM denotes a draw from the stochastic agent-based dynamics, $d_1 ... d_T$ are the time series output by the simulation, and $o_1 ... o_T$ are the number of deaths observed on the corresponding dates. We model the observations as drawn from a negative binomial distribution (appropriate for overdispersed count data) with dispersion parameter σ^2_{obs}. We separately estimated σ^2_{obs} by fitting an autoregressive negative binomial regression to the observed counts using the R package tscount [44]. The negative binomial observation model was strongly preferred to a Poisson model (see Table S2 with AIC values). Together, the likelihood function is given by

$$
\mathcal{L}(p_{\text{inf}}, d_{\text{mult}}, t_0, d_1 ... d_T) = \prod_{t=1}^{t} \Pr\left[o_t | d_t, \sigma^2_{obs}\right].
$$

To obtain the posterior distribution, we use Latin hypercube sampling to draw many (10-80 thousand per location, depending on the size of the prior ranges) samples from the joint uniform prior over $p_{\text{inf}}, d_{\text{mult}}$ and t_0, and then sample the latent variables $d_1 ... d_T$ at each combination of parameters. We compute the likelihood for the full sample (including the latent variables). This allows us to use importance sampling to resample values of $(p_{\text{inf}}, d_{\text{mult}}, t_0, d_1 ... d_T)$ according to the posterior distribution. Finally, we marginalize out $d_1 ... d_T$ to obtain the posterior over the parameters $p_{\text{inf}}, d_{\text{mult}}, t_0$, along with unobservable state variables of the simulation such as the number of infected individuals at each step.

Acknowledgments

This work was supported in part by the Army Research Office by grant MURI W911NF1810208 and in part by grant T32HD040128 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development, National Institutes of Health. Killian was supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1745303. Perrault and Jabbari were supported by the Harvard Center for Research on Computation and Society. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

[1] David Baud, Xiaolong Qi, Karin Nielsen-Saines, Didier Musso, Léo Pomar, and Guillaume Favre. Real estimates of mortality following COVID-19 infection. *The Lancet*, 2020.

[2] Center for Systems Science and Engineering at Johns Hopkins University. Coronavirus COVID-19 Global Cases, 2020. `https://coronavirus.jhu.edu/map.html`.

[3] Fei Zhou, Ting Yu, Ronghui Du, Guohui Fan, Ying Liu, Zhibo Liu, Jie Xiang, Yeming Wang, Bin Song, Xiaoying Gu, et al. Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study. *The Lancet*, 2020.

[4] Bo Xu, Alomía de Gutiérrez, Sumiko Mekaru, Kara Sewalk, Lauren Goodwin, Alyssa Loskill, Emily Cohn, Yulin Hswen, Sarah Hill, María Mercedes Cobo, Alexander Zarebski, et al. Epidemiological data from the COVID-19 outbreak, real-time case information. *Scientific Data*, 7, 2020.

[5] Julien Riou and Christian Althaus. Pattern of early human-to-human transmission of Wuhan 2019 novel coronavirus (2019-nCoV), December 2019 to January 2020. *Eurosurveillance*, 25(4), 2020.

[6] Ruiyun Li, Sen Pei, Bin Chen, Yimeng Song, Tao Zhang, Wan Yang, and Jeffrey Shaman. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV2). *Science*, 2020.

[7] Adam Kucharski, Timothy Russell, Charlie Diamond, Yang Liu, John Edmunds, Sebastian Funk, and Rosalind Eggo. Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The Lancet Infectious Diseases*, 2020.

[8] Kiesha Prem, Alex Cook, and Mark Jit. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLoS Computational Biology*, 13(9):e1005697, 2017.

[9] Robert Verity, Lucy Okell, Ilaria Dorigatti, Peter Winskill, Charles Whittaker, Natsuko Imai, Gina Cuomo-Dannenburg, Hayley Thompson, Patrick Walker, Han Fu, et al. Estimates of the severity of COVID-19 disease. *medRxiv*, 2020.

[10] NYC Department of Health and Mental Hygiene. Coronavirus disease 2019 (covid-19) daily data summary, 2020. `https://www1.nyc.gov/assets/doh/downloads/pdf/imm/covid-19-daily-data-summary-deaths-05172020-1.pdf`.

[11] Graziano Onder, Giovanni Rezza, and Silvio Brusaferro. Case-fatality rate and characteristics of patients dying in relation to covid-19 in italy. *JAMA*, 2020.

[12] Juanjuan Zhang, Maria Litvinova, Yuxia Liang, Yan Wang, Wei Wang, Shanlu Zhao, Qianhui Wu, Stefano Merler, Cecile Viboud, Alessandro Vespignani, et al. Changes in contact patterns shape the dynamics of the covid-19 outbreak in china. *Science*, 2020.

[13] Google. COVID-19 community mobility reports, 2020. `https://www.google.com/covid19/mobility/`.

[14] Josh Katz, Denise Lu, and Margot Sanger-Katz. What is the real coronavirus death toll in each state? *The New York Times*, 2020. `https://www.nytimes.com/interactive/2020/05/05/us/coronavirus-death-toll-us.html`.

[15] Chirag Modi, Vanessa Boehm, Simone Ferraro, George Stein, and Uros Seljak. Total covid-19 mortality in italy: Excess mortality and age dependence through time-series analysis. *medRxiv*, 2020.

[16] Maimuna Majumder and Kenneth Mandl. Early in the epidemic: Impact of preprints on global discourse of 2019-nCOV transmissibility. *SSRN*, 2020.

[17] Giorgio Guzzetta, Piero Poletti, Marco Ajelli, Filippo Trentini, Valentina Marziano, Danilo Cereda, Marcello Tirani, Giulio Diurno, Annalisa Bodina, Antonio Barone, Lucia Crottogini, Maria Gramegna, Alessia Melegaro, and Stefano Merler. Potential short-term outcome of an uncontrolled COVID-19 epidemic in Lombardy, Italy, February to March 2020. *Eurosurveillance*, 25(12):2000293, 2020. doi: https://doi.org/10.2807/1560-7917.ES.2020.25.12.2000293. URL https://www.eurosurveillance.org/content/10.2807/1560-7917.ES.2020.25.12.2000293.

[18] British Broadcasting Corporation. Coronavirus: China outbreak city Wuhan raises death toll by 50%, 2020. https://www.bbc.com/news/world-asia-china-52321529, Last Accessed: 2020-05-17.

[19] CBS New York. Coronavirus Antibodies Present In Nearly 25% Of All NYC Residents, 2020. https://newyork.cbslocal.com/2020/04/27/coronavirus-antibodies-present-in-nearly-25-of-all-nyc-residents/, Last Accessed: 2020-05-17.

[20] Karen Yourish, K.K. Rebecca Lai, Danielle Ivory, and Mitch Smith. One-third of all U.S. coronavirus deaths are nursing home residents or workers. *The New York Times*, 2020. https://www.nytimes.com/interactive/2020/05/09/us/coronavirus-cases-nursing-homes-us.html.

[21] Centers for Disease Control and Prevention. People who are at higher risk for severe illness, 2020. https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/people-at-higher-risk.html.

[22] Peter Mateyka, Melanie Rapino, and Liana Christin Landivar. Home-based workers in the United States, 2012. https://www.census.gov/prod/2012pubs/p70-132.pdf.

[23] U.S. Bureau of Labor Statistics. Labor force statistics from the current population survey, 2019. https://www.bls.gov/cps/cpsaat08.htm.

[24] Pablo De Salazar, Rene Niehus, Aimee Taylor, Caroline Buckee, and Marc Lipsitch. Using predicted imports of 2019-nCoV cases to determine locations that may not be identifying all imported cases. *medRxiv*, 2020.

[25] Stephen Kissler, Christine Tedijanto, Marc Lipsitch, and Yonatan Grad. Social distancing strategies for curbing the COVID-19 epidemic. *medRxiv*, 2020.

[26] Joel Hellewell, Sam Abbott, Amy Gimma, Nikos Bosse, Christopher Jarvis, Timothy Russell, James Munday, Adam Kucharski, John Edmunds, Sebastian Funk, and Rosalind Eggo. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health*, 2020.

[27] Shuo Feng, Chen Shen, Nan Xia, Wei Song, Mengzhen Fan, and Benjamin Cowling. Rational use of face masks in the COVID-19 pandemic. *The Lancet Respiratory Medicine*, 2020.

[28] Gabriella Di Giuseppe, Rossella Abbate, Luciana Albano, Paolo Marinelli, and Italo Angelillo. A survey of knowledge, attitudes and practices towards avian influenza in an adult population of Italy. *BMC Infectious Diseases*, 8(1):36, 2008.

[29] Timothy Russell, Joel Hellewell, Sam Abbott, Christopher Jarvis, Kevin van Zandvoort, Stefan Flasche, Rosalind Eggo, John Edmunds, and Adam Kucharski. Using a delay-adjusted case fatality ratio to estimate under-reporting, 2020. https://cmmid.github.io/topics/covid19/severity/global_cfr_estimates.html. Accessed 03-26-20.

[30] Chinese Center for Disease Control and Prevention. The epidemiological characteristics of an outbreak of 2019 novel coronavirus diseases (COVID-19). *China CDC Weekly*, 2(8):113–122, 2020. http://weekly.chinacdc.cn//article/id/e53946e2-c6c4-41e9-9a9b-fea8db1a8f51.

[31] Mark Parascandola and Lin Xiao. Tobacco and the lung cancer epidemic in china. *Translational Lung Cancer Research*, 8: S21–S30, 2019.

[32] Alessandra Lugo, Piergiorgio Zuccaro, Roberta Pacifici, Giuseppe Gorini, Paolo Colombo, Carlo La Vecchia, and Silvano Gallus. Smoking in italy in 2015-2016: Prevalence, trends, roll-your-own cigarettes, and attitudes towards incoming regulations. *Tumori Journal*, 103(4):353–359, 2017.

[33] Maryellen Guinan, Maryanne McGuckin-Guinan, and Alice Sevareid. Who washes hands after using the bathroom? *American Journal of Infection Control*, 25(5):424–425, 1997.

[34] Durell Johnson, Danielle Sholcosky, Karen Gabello, Robert Ragni, and Nicole Ogonosky. Sex differences in public restroom handwashing behavior associated with visual behavior prompts. *Perceptual and Motor Skills*, 97(3):805–810, 2003.

[35] Rudragouda Channappanavar, Craig Fett, Matthias Mack, Patrick Ten Eyck, David Meyerholz, and Stanley Perlman. Sex-based differences in susceptibility to severe acute respiratory syndrome coronavirus infection. *The Journal of Immunology*, 198(10):4046–4053, 2017.

[36] British Broadcasting Corporation. Coronavirus: South Korea emergency measures as infections increase, 2020. https://www.bbc.com/news/world-asia-51582186.

[37] Pauline Van den Driessche, Michael Li, and James Muldowney. Global stability of SEIRS models in epidemiology. *Canadian Applied Mathematics Quarterly*, 7:409–425, 1999.

[38] Frank Ball, Edward Knock, and Philip O'Neill. Stochastic epidemic models featuring contact tracing with delays. *Mathematical Biosciences*, 266:23–35, 2015.

[39] Gregory Roth, Degu Abate, Kalkidan Hassen Abate, Solomon Abay, Cristiana Abbafati, Nooshin Abbasi, Hedayat Abbastabar, Foad Abd-Allah, Jemal Abdela, Ahmed Abdelalim, et al. Global, regional, and national age-sex-specific mortality for 282 causes of death in 195 countries and territories, 1980–2017: A systematic analysis for the Global Burden of Disease Study 2017. *The Lancet*, 392(10159):1736–1788, 2018.

[40] Paul Allison. *Survival analysis using SAS: a practical guide*. SAS Institute, 2010.

[41] David Collett. *Modelling survival data in medical research*. CRC Press, 2015.

[42] Stephen Lauer, Kyra Grantz, Qifang Bi, Forrest Jones, Qulu Zheng, Hannah Meredith, Andrew Azman, Nicholas Reich, and Justin Lessler. The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: Estimation and application. *Annals of Internal Medicine*, 2020.

[43] Yang Liu, Rosalind Eggo, and Adam Kucharski. Secondary attack rate and superspreading events for SARS-CoV-2. *The Lancet*, 2020.

[44] Tobias Liboschik, Konstantinos Fokianos, and Roland Fried. tscount: An r package for analysis of count time series following generalized linear models. *Journal of Statistical Software*, 2015.

[45] Yan Bai, Lingsheng Yao, Tao Wei, Fei Tian, Dong-Yan Jin, Lijuan Chen, and Meiyun Wang. Presumed asymptomatic carrier transmission of COVID-19. *The Journal of the American Medical Association*, 2020.

[46] Camilla Rothe, Mirjam Schunk, Peter Sothmann, Gisela Bretzel, Guenter Froeschl, Claudia Wallrauch, Thorbjörn Zimmer, Verena Thiel, Christian Janke, Wolfgang Guggemos, et al. Transmission of 2019-nCoV infection from an asymptomatic contact in Germany. *New England Journal of Medicine*, 2020.

[47] Zhanwei Du, Xiaoke Xu, Ye Wu, Lin Wang, Benjamin Cowling, and Lauren Ancel Meyers. Serial interval of COVID-19 among publicly reported confirmed cases. *Emerging Infectious Diseases*, 2020.

[48] Zhan Hu and Xizhe Peng. Household changes in contemporary China: An analysis based on the four recent censuses. *The Journal of Chinese Sociology*, 2(1):9, 2015.

[49] Dan He, Xuying Zhang, Zhili Wang, and Yu Jiang. China fertility report, 2006–2016. *China Population and Development Studies*, 2(4):430–439, 2019.

[50] United Nations. World population prospects 2019. 2019. `https://population.un.org/wpp/`.

[51] Statista. Household structures in Italy in 2018, 2018. `https://www.statista.com/statistics/730604/family-structures-italy/`, Last Accessed: 2020-03-28.

[52] Statista. Number of single-person households in Italy from 2012 to 2018, 2018. `https://www.statista.com/statistics/728061/number-of-single-person-households-italy/`, Last Accessed: 2020-03-28.

[53] Statista. Number of couples with children in Italy from 2012 to 2018, by number of children, 2018. `https://www.statista.com/statistics/570106/number-of-couples-with-children-italy/`, Last Accessed: 2020-03-28.

[54] Statista. Biennial average number of household members in Italy from 2012 to 2018, 2018. `https://www.statista.com/statistics/671945/biennial-average-number-of-families-with-children-italy/`, Last Accessed: 2020-03-28.

[55] Statista. Number of single parents in Italy from 2011 to 2018, by number of children, 2018. `https://www.statista.com/statistics/570234/number-of-single-parents-with-children-in-italy-by-number-of-children/`, Last Accessed: 2020-03-28.

[56] Elisabetta Carrà, Margherita Lanz, and Semira Tagliabue. Transition to adulthood in Italy: An intergenerational perspective. *Journal of Comparative Family Studies*, 45(2):235–248, 2014.

[57] Minnesota Population Center. Integrated public use microdata series, international: Version 7.2 [dataset], 2019. `https://doi.org/10.18128/D020.V7.2`.

[58] Yu Xu, Limin Wang, Jiang He, Yufang Bi, Mian Li, Tiange Wang, Linhong Wang, Yong Jiang, Meng Dai, Jieli Lu, Min Xu, Yichong Li, Nan Hu, Jianhong Li, Shengquan Mi, Chung-Shiuan Chen, et al. Prevalence and Control of Diabetes in Chinese Adults. *The Journal of the American Medical Association*, 310(9):948–959, 09 2013.

[59] Zengwu Wang, Zuo Chen, Linfeng Zhang, Xin Wang, Guang Hao, Zugui Zhang, Lan Shao, Ye Tian, Ying Dong, Congyi Zheng, Jiali Wang, Manlu Zhu, William Weintraub, and Runlin Gao. Status of hypertension in China. *Circulation*, 137(22): 2344–2356, 2018.

[60] Pietro Modesti, Maria Calabrese, Ilaria Marzotti, Hushao Bing, Danilo Malandrino, Maria Boddi, Sergio Castellani, and Dong Zhao. Prevalence, awareness, treatment, and control of hypertension among Chinese first-generation migrants and Italians in Prato, Italy: The CHIP study. *International Journal of Hypertension*, 2017.

[61] Centers for Disease Control and Prevention. New York City diabetes ABC profile 2011–2012, 2017.

[62] New York City Department of Health and Mental Hygiene. Hypertension in New York City: Disparities in prevalence, 2016.

[63] Yukako Tatsumi and Takayoshi Ohkubo. Hypertension with diabetes mellitus: significance from an epidemiological perspective for Japanese. *Hypertension Research*, 40(9):795–806, 2017.

[64] CDC COVID-19 Response Team. Severe outcomes among patients with coronavirus disease 2019 (COVID-19)—United States, February 12–March 16, 2020, 2020. https://www.cdc.gov/mmwr/volumes/69/wr/mm6912e2.htm.

[65] Xizhe Peng. China's demographic history and future challenges. *Science*, 333(6042):581–587, 2011.

[66] South China Morning Post. Coronavirus: China's first confirmed COVID-19 case traced back to November 17, 2020. https://www.scmp.com/news/china/society/article/3074991/coronavirus-chinas-first-confirmed-covid-19-case-traced-back.

[67] Italian National Institute of Statistics. Median age in Lombardy, 2019. https://www4.istat.it/it/lombardia/dati?qt=gettable&dataset=DCIS_INDDEMOG1&dim=21,0,0, Last Accessed: 2020-03-28.

[68] CIA World Factbook. Field listing: median age, 2020. https://www.cia.gov/library/publications/the-world-factbook/fields/343.html, Last Accessed: 2020-03-28.

[69] Lorenzo Tondo. Italy charges more than 40,000 people with violating lockdown. *The Guardian*, 2020. https://www.theguardian.com/world/2020/mar/18/italy-charges-more-than-40000-people-violating-lockdown-coronavirus.

[70] Fabrizio Carinci. COVID-19: Preparedness, decentralisation, and the hunt for patient zero. *British Medical Journal*, 2020.

[71] Politico. Italian doctors on coronavirus frontline face tough calls on whom to save, 2020. https://www.politico.eu/article/coronavirus-italy-doctors-tough-calls-survival/.

[72] Yascha Mounk. The extraordinary decisions facing Italian doctors, 2020. https://www.theatlantic.com/ideas/archive/2020/03/who-gets-hospital-bed/607807/.

[73] Ana S Gonzalez-Reiche, Matthew M Hernandez, Mitchell Sullivan, Brianne Ciferri, Hala Alshammary, Ajay Obla, Shelcie Fabre, Giulio Kleiner, Jose Polanco, Zenab Khan, et al. Introductions and early spread of sars-cov-2 in the new york city area. *medRxiv*, 2020.

[74] Benedict Carey and James Glanz. Hidden outbreaks spread through U.S. cities far earlier than americans knew, estimates say. *The New York Times*, 2020. https://www.nytimes.com/2020/04/23/us/coronavirus-early-outbreaks-cities.html.

[75] U.S. Bureau of Labor Statistics. American time use survey, 2018. https://www.bls.gov/tus/.

[76] Noreen Qualls, Alexandra Levitt, Neha Kanade, Narue Wright-Jegede, Stephanie Dopson, Matthew Biggerstaff, Carrie Reed, Amra Uzicanin, CDC Community Mitigation Guidelines Working Group Group, et al. Community mitigation guidelines to prevent pandemic influenza—united states, 2017. *CDC MMWR Recommendations and Reports*, 66(1):1, 2017.

[77] Oklahoma County Department of Health. Social distancing fact sheet, 2019. https://www.occhd.org/application/files/3715/7013/9751/Social_Distancing.pdf, Last Accessed: 2020-04-05.

[78] Faruque Ahmed, Nicole Zviedrite, and Amra Uzicanin. Effectiveness of workplace social distancing measures in reducing influenza transmission: a systematic review. *BMC Public Health*, 18:518, 2018.

[79] Peter Totterdell. Work schedules. In Julian Barling, E Kevin Kelloway, and Michael R Frone, editors, *Handbook of work stress*, chapter 3. Sage publications, 2004.

[80] LeaAnne DeRigne, Patricia Stoddard-Dare, and Linda Quinn. Workers without paid sick leave less likely to take time off for illness or injury compared to those with paid sick leave. *Health Affairs*, 35(3):520–527, 2016.

[81] Min W. Fong, Huizhi Gao, Jessica Y. Wong, Jingyi Xiao, Eunice Y.C. Shiu, Sukhyun Ryu, and Benjamin J. Cowling. Nonpharmaceutical measures for pandemic influenza in nonhealthcare settings—social distancing measures. *Emerging Infectious Diseases*, 26(5), 2020.

[82] Francesco Aimone. The 1918 influenza epidemic in new york city: a review of the public health response. *Public Health Reports*, 125(3_suppl):71–79, 2010.

[83] Liisa Ecola and Thomas Light. Equity and congestion pricing. *RAND Corporation*, pages 1–45, 2009.

[84] NBC News. Walmart will limit customers and create one-way traffic inside its stores, 2020. https://www.nbcnews.com/news/us-news/walmart-will-limit-customers-create-one-way-traffic-inside-its-n1176461.

[85] NBC Connecticut. New social distancing measures take effect at stores across Connecticut, 2020. https://www.nbcconnecticut.com/news/local/new-social-distancing-measures-enacted-at-stores-across-connecticut/2249937/.

[86] Alina Selyukh. Supermarkets add 'senior hours' for vulnerable shoppers, 2020. https://www.npr.org/sections/coronavirus-live-updates/2020/03/19/818488098/supermarkets-add-senior-hours-for-vulnerable-shoppers.

[87] Centers for Disease Control and Prevention. Interim pre-pandemic planning guidance: community strategy for pandemic influenza mitigation in the united states. 2007.

[88] WGNTV. This student created a network of 'shopping angels' to help the elderly get groceries during the COVID-19 pandemic, 2020. https://wgntv.com/news/this-student-created-a-network-of-shopping-angels-to-help-the-elderly-get-groceries-during-the-covid-

[89] American Broadcasting Company. Santa Monica volunteers protect seniors with new grocery delivery service, 2020. https://abc7.com/coronavirus-seniors-covid-19-groceries/6045815/.

[90] Kathy Gardner and Kate Lister. 2017 state of telecommuting in the U.S. employee workforce, 2017. https://www.flexjobs.com/2017-State-of-Telecommuting-US/.

[91] Timothy Noah. 'it makes me very angry': Coronavirus damage ripples across the workforce, 2020. https://www.politico.com/news/2020/03/09/america-workers-outbreak-uncertainty-124191.

[92] Drew Desilver. As coronavirus spreads, which U.S. workers have paid sick leave – and which don't?, 2020. https://www.pewresearch.org/fact-tank/2020/03/12/as-coronavirus-spreads-which-u-s-workers-have-paid-sick-leave-and-which-dont/.

[93] David Card. Labor supply with a minimum hours threshold. In *Carnegie-Rochester Conference Series on Public Policy*, volume 33, pages 137–168. Elsevier, 1990.

[94] Centers for Disease Control and Prevention. Interim pre-pandemic planning guidance: Community strategy for pandemic influenza mitigation in the united states— early, targeted, layered use of nonpharmaceutical interventions, 2007. https://www.cdc.gov/flu/pandemic-resources/pdf/community_mitigation-sm.pdf.

[95] World Health Organization China. Report of the WHO-China Joint Mission on Coronavirus Disease 2019 (COVID-19), 2020. https://www.who.int/docs/default-source/coronaviruse/who-china-joint-mission-on-covid-19-final-report.pdf.

[96] Qun Li, Xuhua Guan, Peng Wu, Xiaoye Wang, Lei Zhou, Yeqing Tong, Ruiqi Ren, Kathy Leung, Eric Lau, Jessica Wong, et al. Early transmission dynamics in Wuhan, China, of novel coronavirus–infected pneumonia. *New England Journal of Medicine*, 2020.

Appendix

Methods

Model description

We develop an agent-based model for COVID-19 spread which accounts for the distributions of age, household types, comorbidities, and contact between different age groups in a given population. The model follows a *susceptible-exposed-infectious-removed (SEIR)* template [37, 38].

Specifically, we simulate a population of n agents (or individuals), each with an age a_i, a set of comorbidities c_i, and a household (a set of other agents). We stratify age into ten-year intervals and incorporate hypertension and diabetes as comorbidities. These comorbidities are common worldwide [39] and have been associated with a higher risk of in-hospital death for COVID-19 patients [3]. However, our model can be expanded to include other comorbidities of interest in the future. The specific procedure we use to sample agents from the joint distribution of age, household structures, and comorbidities is described below.

The simulation tracks two states for each individual: the *infection state* and the *isolation state*. The infection state is divided into $\{$*susceptible, exposed, infectious, removed*$\}$. *Susceptible* individuals are those who have never been contacted by an infectious individual. *Exposed* individuals are those who have had contact with an infectious individual, though not all exposed individuals become infectious. If an exposed individual contracts the disease, they proceed to the infectious state.[2] *Infectious* is further

[2]Currently, our simulation implementation does not separately track individuals who are exposed but do not become infected, and instead groups them with the susceptible population. This is because we assume that, if exposed again, they will become infected with the same

subdivided into severity levels {*asymptomatic, mild, severe, critical*}. We interpret mild severity as symptomatic (but not requiring hospitalization), severe as requiring hospitalization, and critical as eligible for intensive care unit (ICU) care. The *removed* state is further subdivided into {*recovered*, *deceased*}. Individuals in all severity levels can transmit the disease, but those in the *asymptomatic* state do so at a rate $\alpha < 1$ times that of symptomatic cases. The decision to incorporate reduced transmission for asymptomatic individuals is based on the fact that, though infection by asymptomatic individuals has been observed in case clusters and in examinations of serial intervals [45–47], available evidence suggests that individuals with no or limited symptoms are less infectious than those with severe symptoms [6]. Currently, our simulation incorporates two levels of infectiousness (before and after the onset of symptoms), but it can be adjusted as better information on how viral shedding increases with severity of illness becomes available. We acknowledge that our assumptions surrounding transmissibility and disease severity – as derived from existing literature – may serve as a limitation of our model, as many of these factors are evolving over time.

Each individual has a separate isolation state {*isolated*, *not isolated*}. If isolated, the individual is unable to infect others. We assume that (1) asymptomatic individuals are never isolated, (2) mild individuals become isolated over a mean time of λ_{isolate} days (see Table 1) after the onset of symptoms, and (3) all severe and critical individuals are isolated. However, our simulation framework can easily accommodate different sets of assumptions about isolation (for example, preemptively isolating exposed individuals if they are known to have had contact with an infectious agent).

The disease is transmitted over a contact structure, which is divided into in-household and out-of-household groups. Each agent has a household consisting of a set of other agents. Individuals infect members of their households at a higher rate than out-of-household agents. We model out-of-household transmission using country-specific estimated contact matrices [8]. These matrices state the mean number of daily contacts an individual of a particular age strata has with individuals from each of the other age strata. We assume demographics (including age and household distribution) in Hubei and Lombardy are well-approximated by country-level data.

The model iterates over a series of discrete time steps, each representing a single day, from a starting time t_0 to an end time T. There are two main components to each time step: disease progression and new infections. The progression component is modeled by drawing two random variables for each individual each time they change severity levels (e.g., on entering the mild state). The first random variable is Bernoulli and indicates whether the individual will recover or progress to the next severity level. The second variable represents the amount of time until progression to the next severity level. We use exponential distributions for almost all time-to-event distributions, a common choice in the absence of specific distributional information [40, 41]. The exception is the incubation time between asymptomatic and mild states, where more specific information is available; here, we use a log-normal distribution (see $\mu_{e\to m}$ and $\sigma^2_{e\to m}$ in Table 1) based on estimates by [42]. Table 1 summarizes all distributions and their parameters.

In the new infections component, individuals in the susceptible state may enter the exposed state. Infected individuals infect each of their household members with probability p_h at each time step. p_h is calibrated so that the total probability of infecting a household member before either isolation or recovery matches the estimated secondary attack rate for household members of COVID-19 patients (i.e., the average fraction of household members infected) [43]. Infected individuals draw outside-of-household contacts from the general population using the country-specific contact matrix. For an infected individual of age group i, we sample $w^s_{ij} \sim \text{Poisson}(M^s_{ij})$ contacts for each age group j and setting s where M^s is the country-specific contact matrix for setting s. We include contacts in work, school, and community settings. Poisson distributions are a standard choice for modeling contact distributions [8]. Then, we sample w^s_{ij} contacts of age j uniformly with replacement, and each contact is infected with the probability p_{inf}, the probability of infection given contact. There is evidence to suggest that the probability of infection is higher for an older individual than younger given the same exposure [12], consistent with decline in immune function with age. We adjust for this by letting the probability of infection be βp_{inf} when the exposed individual is over the age of 60, for $\beta > 1$. β is calibrated to match the fraction of deaths in China attributed to individuals over the age of 60, resulting in a value of 1.25. This is consistent with the relationship between age and attack rate amongst close contacts of a confirmed case reported by [12], where the increase in risk of infection for a contact over 65 years old was estimated in the range 1.12–1.92.

Sampling agents

Our process for sampling agents follows three steps that successively sample households, individual agents within households, and comorbidities for each agent. Because the full joint distributions over all of these quantities are not known, we implement a sampling procedure that respects the marginal distributions of household structure and age, as well as the marginal distribution for the occurrence of comorbidities within each age group.

First, we use information on the distribution of household structures to draw a type of household (e.g., single person, couple, nuclear family, or multigenerational family). Second, we sample the ages of the individual agents according to their role in

probability as an individual who has never been exposed. However, the implementation can be modified to support either differing probabilities of contracting the disease after first exposure or policies that treat exposed and susceptible individuals differently.

the household (e.g., parent, child, or grandparent) combined with information about the age distribution of the population and the intergenerational interval. For China, we use household distributions from the 2010 Chinese census [48], intergenerational intervals from [49], and the age distribution provided by UN population statistics [50]. For Italy, we use demographic statistics from Statista online portal about the following: household structure distribution [51], single-person households [52], couples with children [53] and corresponding family size [54], and single parents with children [55]. Furthermore, we assume that children could stay within the family until the age of 30 and that couples without children were aged 30+, to account for societal patterns reported in familial studies which may have affected household distribution metrics [56]. In New York City, we circumvent these steps by instead sampling individual households directly from census microdata. We use the public use microdata from the 2015 American Community Survey [57]. We draw from household-level responses located in New York City, repeatedly sampling a household of individuals with their reported ages until the desired population size (8.4 million) is reached.

Third, we sample comorbidities from the corresponding country- and age-specific distributions. For China, we use estimates on age-specific prevalence of diabetes [58] and hypertension [59]. For Italy, we use estimates from the Global Burden of Disease study on diabetes [39] and a recent study of age-stratified hypertension prevalence [60]. For New York City, we use city-level estimates of age-specific prevalence for both comorbidites [61, 62]. We ensure that diabetes and hypertension are appropriately correlated using a single global estimate for the probability of hypertension in individuals with diabetes [63].

Estimating disease progression from age and comorbidities

Many of the parameters for this model are assigned values based on estimates in the literature, shown in Table 1. However, we currently lack a detailed understanding of the joint impact of age and comorbidities on disease progression and mortality. Currently, estimated infection fatality rates (IFRs) are available by age but not for each specific combination of age and comorbidities. To obtain these specific estimates, we model the IFR with a logistic regression fit to IFRs estimated by Verity et al. [9] on data from mainland China. The logistic model is discussed in the next section. This model yields $p_{m\to d}(a_i, c_i)$, the country-independent probability that an individual i of age a_i and comborbidity status c_i will die if infected with SARS-CoV-2. Corrections for country-specific differences in mortality are handled via the parameter d_{mult}.

The simulation also requires specific values for the probabilities of transitioning between the disease states mild, severe, critical, and death. However, there is currently insufficient information available to infer the probabilities of these individual transitions for each combination of age and comorbidity. We assume that while the absolute values of these probabilities may vary based on age and comorbidity, the *ratios* between them do not exhibit such strong dependency. In particular, we assume that there are coefficients $\gamma_{s\to c}(a_i)$ and $\gamma_{c\to d}$ such that $p_{s\to c}(a_i, c_i) = \gamma_{s\to c}(a_i)p_{m\to s}(a_i, c_i)$ and $p_{c\to d}(a_i) = \gamma_{c\to d}p_{m\to s}(a_i, c_i)$. We allow $\gamma_{s\to c}(a_i)$ to be age-specific while assuming that $\gamma_{c\to d}$ is age-homogeneous because of the information currently available to estimate them; namely, we set $\gamma_{s\to c}(a_i)$ based on the estimated probabilities of hospitalization from [9] and ICU admission by age group in the US from [64] and $\gamma_{c\to d}$ based on the probability of death for all critical patients in China [30]. Note that we assume both coefficients to be independent of the comorbodities c_i. Then, we can solve for $p_{m\to s}(a_i, c_i)$ such that

$$p_{m\to s}(a_i, c_i) \cdot \gamma_{s\to c}(a_i)p_{m\to s}(a_i, c_i) \cdot \gamma_{c\to d}p_{m\to s}(a_i, c_i) = p_{m\to d}(a_i, c_i),$$

and set $p_{s\to c}(a_i, c_i)$ and $p_{c\to d}(a_i, c_i)$ accordingly. Future work can relax the assumptions in this process as more information becomes available about how age and comorbidity impact the progression between disease states.

Estimating mortality from age and comorbidities

We require a model of $p_{m\to d}(a_i, c_i)$, however existing data sources only specify $p_{m\to d}(a_i)$ and $p_{m\to d}(c_i)$. To infer the joint distribution, we assume a linear (logistic) interaction between age bracket, diabetes status, and hypertension status. Specifically, we assume

$$p_{m\to d}(a_i, c_i) = \sigma\Big(\beta_{\text{age}}(a_i) + \beta_{\text{diabetes}}\,\mathbb{1}\left[\text{diabetes} \in c_i\right] + \beta_{\text{hypertension}}\,\mathbb{1}\left[\text{hypertension} \in c_i\right]\Big),$$

where $\beta_{\text{age}}(a_i)$ has a value for each age bracket (e.g., 20-30, 30-40, etc., 7 in total) and β_{diabetes} and $\beta_{\text{hypertension}}$ are scalars.

The marginal distribution $p_{m\to d}(a_i)$ is taken from [9], which corrects for underreporting of infections in China. To obtain a comparable marginal distribution $p_{m\to d}(c_i)$, we scaled the reported CFR for each comorbidity group [30] by an age-adjusted correction for reporting obtained based on [9] (making the assumption that the probability of documentation is independent of comorbidity status after conditioning on age). We obtained data from the literature on the prevalence of diabetes and hypertension [59] in China by age [58], as well as a single global estimate of $p(\text{hypertension}|\text{diabetes})$ [63]. We assume that these distributions are the same in COVID-19 patients as in the general population. Given this information, we use gradient descent to find a set of parameters β which minimize the mean squared error in the following marginal consistency constraints:

$$p_{m\to d}(a_i) = \sum_{\text{diabetes,hypertension}} p(\text{diabetes, hypertension}|a_i) p_{m\to d}(a_i, \text{diabetes, hypertension}), \quad \forall a_i,$$

$$p_{m\to d}(\text{diabetes}) = \sum_{a_i} p(a_i|\text{diabetes}) \sum_{\text{hypertension}} p(\text{hypertension}|a_i, \text{diabetes}) p_{m\to d}(a_i, \text{diabetes, hypertension}),$$

$$p_{m\to d}(\text{hypertension}) = \sum_{a_i} p(a_i|\text{hypertension}) \sum_{\text{diabetes}} p(\text{diabetes}|a_i, \text{hypertension}) p_{m\to d}(a_i, \text{diabetes, hypertension})$$

The set of estimated parameters are

$$\begin{aligned}
&\beta_{\text{age}}(18-30) = -8.49,\\
&\beta_{\text{age}}(30-40) = -7.68,\\
&\beta_{\text{age}}(40-50) = -7.41,\\
&\beta_{\text{age}}(50-60) = -6.39,\\
&\beta_{\text{age}}(60-70) = -5.41,\\
&\beta_{\text{age}}(70-80) = -4.54,\\
&\beta_{\text{age}}(80-100) = -4.05,\\
&\beta_{\text{diabetes}} = 1.22,\\
&\beta_{\text{hypertension}} = 1.58.
\end{aligned}$$

The coefficients should be interpreted relative to the baseline -8.49 value for the 18-30 group. For example, the value -7.68 for the 30-40 group indicates that the log-probability of mortality increases by 0.80 when age is increased from 18-30 to 30-40, holding comordibity status equal. Over 10 random restarts, the marginal values were always fit to within numerical tolerance by the same set of parameters (less than 0.1% maximum difference in the value of a parameter between runs). This suggests that the model parameters are fully identifiable in this setting.

Experimental settings

Experimental settings for Hubei

We draw a population of individuals from the age, household, and comorbidity distributions for China since more specific information is not available for Hubei (though the fraction of individuals over 65 is within the typical range for many Chinese provinces [65]). We simulate a population of 10 million individuals for computational tractability. Note that Hubei has a population of approximately 58 million. However, 10 million is comparable to the population of Wuhan, where most cases were concentrated [30]. After the lockdown, all contact frequencies are reduced by a factor of 99.3, set to obtain the number of outside-of-household contacts reported in post-lockdown surveys [12]. We also modeled closure of schools on the lockdown date.

We set the range of the uniform prior distributions as follows. The prior over p_{inf} was set to contain all values with significant likelihood, with the final range being $[0.020, 0.035]$. The prior over t_0 was set to contain up to 7 days before the first reported case on November 17 [66], and 3 days afterwards (for a set of 10 days total). It is possible that substantial new backdating of the start of the epidemic could alter our results. Finally, the parameter d_{mult} captures variation in IFR, which is not precisely known in any location. We start from age-stratified IFR estimates by Verity et al. [9]. In our model, these values result in an overall IFR of approximately 0.4% (lower than the 0.66% estimated by [9] because attack rates in our model are higher in younger groups, due to the larger numbers of daily contacts in younger groups vs older [8]). We placed a uniform prior over d_{mult} in the range [1,3], and then conditioned in the posterior on the IFR lying in the range 0.4–0.8%. Together, this procedure is designed to allow variation by approximately 50% around the IFR estimated by [9].

Experimental settings for Lombardy

We simulate a population of 10 million individuals (representing the population of Lombardy) drawn from the Italian distribution of age, household structure, and comorbidity status. The full demographic information needed to parameterize the simulation was not available for Lombardy specifically, but available information suggests broadly similar characteristics (e.g., the median age in Lombardy is 45 [67], comparable to Italy in general at 46.5 [68]). After the lockdown on March 8, the number of contacts

for all age groups is reduced by a factor of 10, reflecting a less severe lockdown than in Hubei [69], and consistent with reported declines in mobility associated with retail and recreation in the range 80-95% [13]. In practical terms, such a level of physical distancing induces an average of approximately 1.3 to 1.8 outside of household contacts a day for people aged 15-29 or 30-49, while it corresponds with about 0.75 to 0.9 interactions on average for the 50-69 age group [8]. We also model closure of all schools on the lockdown date.

As in Hubei, we set the prior range for p_{inf} to include all values with significant likelihood, resulting in an interval [0.025, 0.04]. Also as in Hubei, the prior for t_0 was set to be uniform over a range of dates including up to 7 days before the infected travelers reportedly landed in Milan on January 23 [70], and up to 3 days afterwards. We adjusted the way the parameter d_{mult} was applied to account for the substantially different age composition of deaths in Italy than in either Hubei or New York City. Specifically, approximately 95% of reported deaths in Italy were among individuals 60 years or older, compared to approximately 80% in China [30] or 73.6% in New York City [10]. One potential factor which could contribute to disparities in death rates are reports that older individuals in severe condition may have been less likely to receive care under the triage strategies adopted in response to overburdened health systems in Italy [71, 72]. Accordingly, instead of scaling fatality rates uniformly across age groups, we calibrated a multiplier for the fatality rate in the 60+ age group to match the fraction of deaths attributed to that group.

Experimental settings for New York City

We simulate a population of 8.4 million individuals (representing the population of New York City), sampled in household units from census microdata for New York City. We model a two-step reduction in contact, consistent with mobility data [13]. The official lockdown was instituted on March 23, and we model a 10-fold reduction in contact on that date (as in Lombardy). However, mobility data shows that significant reductions in mobility began the week before the official lockdown, suggesting preemptive distancing measures by individuals in anticipation of official policy. Accordingly, we model contact in all non-household settings as reduced by a factor of 2 starting on March 16.

We set the prior range for p_{inf} to be [0.03, 0.07], again set to include all values with non-negligible likelihood. d_{mult} was given a uniform prior over the range [1, 4], allowing for but not mandating a higher IFR than Hubei. We handled the starting conditions of the epidemic differently in New York City than in Lombardy due to reports of multiple distinct importation events over the course of February [73], with modeling studies suggesting the potential for thousands of cases present by the start of March [74]. Instead of attempting to explicitly model multiple importations, we fixed t_0 at February 10 and placed a uniform prior over the number of infected individuals present on that date, in the range [5, 500].

Experimental settings for containment policies

In either scenario, we simulate 50% or 100% of a single age cohort as engaged in salutary sheltering. We model salutary sheltering as removing all between-household contacts for the "sheltered" individuals, though they are still able to infect and be infected by their household members. We model physical distancing as reducing the expected number of daily contacts between non-sheltered individuals of any two age groups by a factor of two. It is worth noting that while there is no difference on average between 100% of the whole population reducing their physical interactions by a factor of 2 (i.e., physical distancing) and having 50% of the whole population engage in salutary sheltering, the variance in the number of daily contacts–limited to household members–an individual may have in the latter is much lower. In each location, we simulate sheltering (and physical distancing if included in the scenario) as starting on the first day of modeled contact reductions for that location. All values are taken in expectation over 1000 values of the model parameters sampled from the posterior distribution for each location.

Mechanisms for physical distancing

Our analysis in Hubei, Lombardy, and New York City suggests that combining salutary sheltering by a portion of the population with physical distancing by the rest can be effective at mitigating the epidemic. In simulation, we modeled physical distancing as reducing the expected number of daily out-of-household contacts between any two members of the population (who are not sheltering) by a factor of two. Here, we consider specific ways that this kind of physical distancing could be achieved by governments, businesses, and community organizations. We consider categories of daily activities defined by the American Time Use Survey [75] and draw on the literature around pandemic preparedness to provide concrete ways that the number of daily contacts can be reduced within each category. We note that this survey draws primarily on American sources. While many suggestions will be broadly applicable, similar sets of recommendations could be compiled specific to particular locations.

Work Businesses can take many actions to reduce the number of workers who are present at the workplace at a given time and to increase the amount of physical distance between workers who are present. Beyond the extreme solution of extended leaves of absence which could hurt the economy in the long run, more viable examples include have workers telecommute (even partially if their job cannot be performed fully remotely), staggering shifts to reduce the number of workers simultaneously

present, increasing physical spacing between workers in an office, limiting the use of shared workstations, offering alternate locations to eat meals instead of a shared lunchroom, and staggering breaks and lunch hours [76, 77]. “Cohorting" – keeping smaller groups of workers together over time so that contact remains mostly within the cohort – may also be helpful at reducing the total number of other people a given individual comes into contact with [76]. [78] surveyed three epidemiological and twelve influenza-related modeling studies, showing that workplace physical distancing measures reduced the cumulative attack rate in the general population. They also established that epidemic peaks were delayed and reduced and that these positive effects were more pronounced when workplace physical distancing measures were coupled with other types of non-pharmaceutical interventions, as well as with therapeutic solutions. Means of instituting shift work, including staggering hours and compressed work-weeks, are reviewed by Totterdell [79], along with evidence about the health and wellness impacts of shift schedules on workers. Finally, an important consideration is availability of sick or personal leave so that symptomatic workers are able to isolate instead of continuing to expose coworkers [80]. As [81] point out, and similar to school closure policies, both the timing and duration of workplace interventions would be critical to affect the course of the pandemic.

Commuting Adoption of shift work by businesses also has the potential to reduce physical contact on public transit. Apart from the scheduling of individual workers, staggering of business hours across a city – such that different businesses open and close at different times – was a tool used by the City of New York during the 1918 influenza pandemic to reduce congestion on public transit [82]. Congestion pricing is another potential tool to reduce peak occupancy on public transit. It is worth noting that congestion pricing may entail equity concerns, which could be mitigated by charging different rates depending on income [83] or by implementing the program instead as a subsidy for using transit during off-peak hours.

Grocery shopping Several measures have been adopted by grocery stores during the COVID-19 pandemic to reduce the number of shoppers who come into contact with each other. These include making aisles one-way [84], limiting the number of individuals in the store at a given time [84], marking appropriate spacing on the floor in checkout lines [85], and reserving specific hours for members of a particular age group (e.g., older groups) [86]. CDC pandemic preparedness guidelines also recommend that individuals reduce their frequency of trips to the grocery store [87]. This could be instituted either on an individual/voluntary basis or via governments or stores assigning specific time periods for individuals or groups of individuals to shop. Some communities have also used volunteers to delivery groceries to individuals who are older or otherwise more vulnerable [88, 89].

Sports, exercise, and recreation Gyms could stagger the times that individuals use the facility and take steps to limit close contacts, e.g., ensuring appropriate spacing between people in a group exercise class. Improving ventilation in gyms may also be important to avoid increased transmission rates due to physical exertion in an enclosed space.

Attending or hosting social events Restaurants and bars could be required to maintain appropriate spacing between parties. Events involving more than a specific number of people could still be restricted.

Religious and spiritual activities In order to reduce the density of contacts at religious services, possible measures include increasing the number of services to reduce the number of people attending each service, cancelling associated social gatherings like coffee hours or youth groups, sending religious school lessons home with students instead of holding sessions in-person, broadcasting services where possible, and providing counseling services by phone instead of in person [77].

Feasibility of physical distancing

The feasibility of instituting proposed mechanisms for physical distancing may be subject to a range of considerations, including economic circumstances. This is particularly relevant for workplace physical distancing measures. We survey available information on several considerations related to the economic feasibility of physical distancing in different populations: telecommuting patterns, availability of sick leave, and consequences of reduced operations by a workplace.

Telecommuting

We consider the most recent statistics as provided by the 2017 State of Telecommuting Report about the U.S. employee workforce [90] to better understand the demographics of home-based workers. Some subsets of the adult population – depending on the age category, occupation type, and industry – were already more likely to work from home than others before this pandemics. Taking these factors into consideration helps rate the feasibility of our proposed policies and provides guidance for immediate intervention.

Age Half of telecommuters in the US are 45 years of age or older, as compared to just 41% of the overall workforce [90]. Specifically, the greatest disproportional participation in home-based work is observed for the following age groups: 65+ years old (odds ratio, OR=1.7) and 55-64 years old (OR=1.2).

Occupation type Among telecommuting jobs, sales, management, office, and administrative occupations are the most common. Together, they account for 43% of home-based jobs, while representing only 34% of non-telecommuting occupations. Five type of occupations are more prevalent among home-based workers as opposed to non-telecommuting workforce participants [90]: computer and mathematical work (OR=2.8), military work (OR=2.5), arts/design/entertainment/sports/media (OR=2.0), personal care and service (OR=2.0), and business and financial operations (OR=1.9).

Industry The largest share of telecommuters is due to three industries [90]: professional, scientific, and technical services (17%); healthcare and social assistance (11.6%); and finance and insurance (9.7%). The prevalence in telecommuting vs. non-telecommuting work is higher in management (OR=32.8), agriculture (OR=3.4), real estate (OR=3), information (OR=3), and mining (OR=2.4).

Access to telecommuting options In the US, having the possibility to telecommute is mostly a function of both the worker's company size and employment status. Telecommuting is more prevalent among large vs. small companies. 12% of firms with more than 500 employees offer such a possibility, whereas only 5% of those with less than 100 employees do so. Moreover, employers are more likely to offer the option to work from home to their full-time rather than part-time employees (8% vs. 2% respectively as of 2017). It is expected that these tendencies would be confirmed in other industrialized countries as well.

Availability of sick leave

Policies related to sick leave vary widely by location. For example, in the United States, workers are not generally guaranteed paid sick leave; only 11 states require employers to offer workers paid leave [91]. As a consequence, 24% have no paid sick leave at all, with strong heterogeneities in access to paid sick leave by income level, occupation type, and activity sector [92]. Notably, 93% of workers in the top tenth of the income distribution receive paid sick leave, compared with only 30% of those in the bottom tenth [91]. However, the adoption of isolation of potentially infected individuals is strongly contingent on access to paid sick or personal leave.

Consequences of reduced operations

For work that cannot be performed remotely, reduced operations due to sheltering or physical distancing interventions will inevitably reduce employment. For a less than total shutdown, work in the labor economics literature considers strategies to allocate the reduced number of hours available to workers that have not been affected by lay-offs, including the use of a minimum hours threshold corresponding with guaranteed earnings [93], as well as a maximum number of hours worked per person to preclude additional employment inequalities. Regardless, additional options to offset the income loss for eligible workers include provisions for unemployment insurance [94].

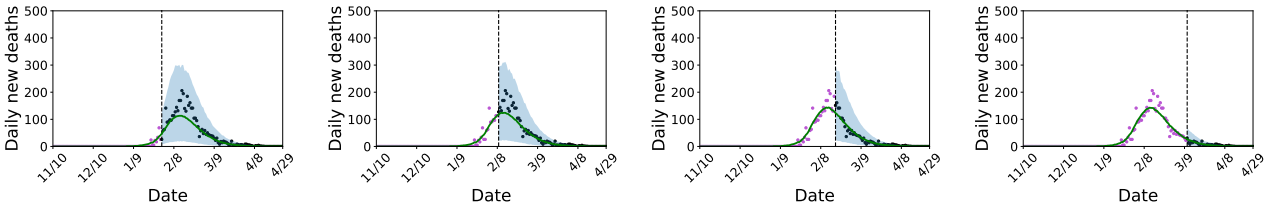

Figure S1: Predictive posterior for Hubei as a function of when the training period ends. Black dashed line: end of training period. Green line: posterior median. Blue shaded region: 90% credible interval. Pink dots: training data. Black dots: held-out data. The 90% credible interval of the predictive posterior includes contains the held-out data at all points, including when the model is fit using only data from the earliest portion of the epidemic.

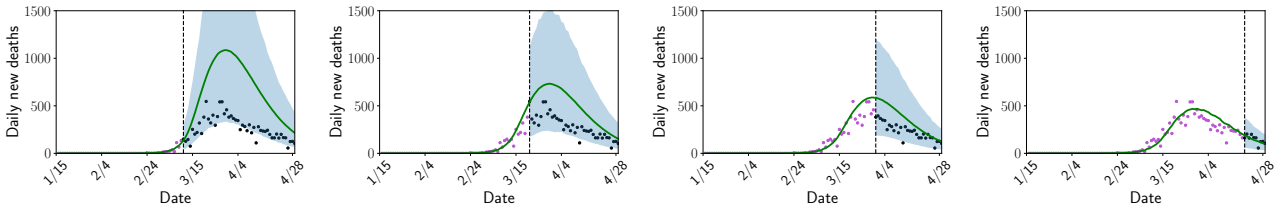

Figure S2: Predictive posterior for Lombardy as a function of when the training period ends. Black dashed line: end of training period. Green line: posterior median. Blue shaded region: 90% credible interval. Pink dots: training data. Black dots: held-out data. The 90% credible interval of the predictive posterior includes contains the held-out data at all almost all points, including when the model is fit using only data from the earliest portion of the epidemic. The model over-predicts deaths early in the epidemic, though the timing of the peak is correctly captured early on. Much of the over-prediction is corrected with additional training data even before the peak is observed.

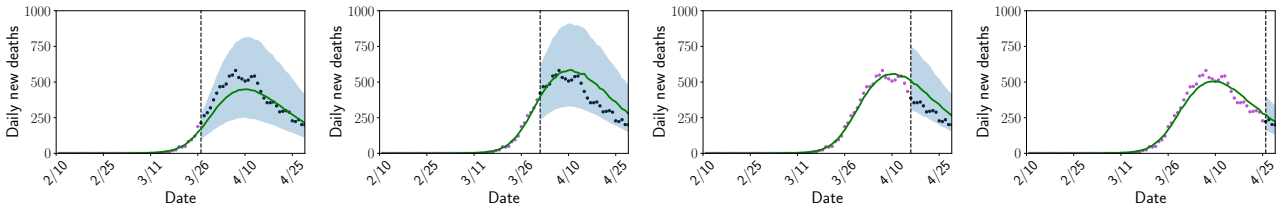

Figure S3: Predictive posterior for New York City as a function of when the training period ends. Black dashed line: end of training period. Green line: posterior median. Blue shaded region: 90% credible interval. Pink dots: training data. Black dots: held-out data. The 90% credible interval of the predictive posterior includes contains the held-out data at all points, including when the model is fit using only data from the earliest portion of the epidemic. Using data from only the earliest stage, the model slightly misidentifies the timing and magnitude of the peak, but these aspects of the prediction substantially improve even without observing the peak in the training data (c.f. the first vs second figure from the left).

Table S1: Model parameters

Parameter	Description	Value and/or source
$p_{m\to s}(a_i, c_i)$	Prob. of progressing from mild to severe given age a_i and comorbidities c_i	Estimated from [9] (see above)
$p_{s\to c}(a_i, c_i)$	Prob. of progressing from severe to critical given age a_i and comorbidities c_i	As above
$p_{c\to d}(a_i, c_i)$	Prob. of progressing from critical to death given age a_i and comorbidities c_i	As above
p_h	Prob. of infecting each household member each day	Calibrated to match [43]
p_{inf}	Prob. of infecting an outside household contact	Free parameter
$\mu_{e\to m}$	Log-mean time to progress from exposed to mild (mean incubation period)	1.621 [42]
$\sigma^2_{e\to m}$	Log-standard deviation time to progress from exposed to mild	0.418 [42]
$\lambda_{m\to s}$	Mean time to progress from mild to severe	7 days [95]
$\lambda_{s\to c}$	Mean time to progress from severe to critical	7.5 days (using 14.5 days from onset to mechanical ventilation in [3])
$\lambda_{c\to d}$	Mean time to progress from critical to death	4.5 days (subtracting $\lambda_{m\to s}$ and $\lambda_{s\to c}$ from onset-to-death in [3])
λ_{isolate}	Mean time for an individual in the mild state to isolate	4.6 days (time to first medical care [96])
$\lambda_{m\to r}$	Mean time to recovery for an individual in the mild state	14 days [95]
$\lambda_{s\to r}$	Mean time to recovery for an individual in the severe state	$28 - \lambda_{m\to s}$ (midpoint of onset-to-recovery for severe [95])
$\lambda_{c\to r}$	Mean time to recovery for an individual in the critical state	$35 - \lambda_{m\to s} - \lambda_{s\to c}$ (midpoint of [95] onset-to-recovery for critical)
α	Reduction in infectiousness before symptoms	0.55 [6] [2]
M	Contact matrix (for each country)	[8]
t_0	First date with at least 5 infected individuals	Free parameter

[2] This setting for α is likely pessimistic in that Li et al.'s estimate for reduction in transmissibility is for undocumented cases, including both asymptomatic cases and those with limited symptoms [6]. Future work should examine the impact of a potentially lower α as better information on transmissibility in the asymptomatic state becomes available.

Table S2: Comparison of Poisson and negative binomial observation models in each location, along with estimated dispersion parameter σ^2_{obs} for the negative binomial. The negative binomial model is strongly preferred by AIC in each location. S2

Location	Poisson AIC	Negative binomial AIC	σ^2_{obs}
Hubei	891.30	670.26	0.337
Lombardy	4741.82	877.97	0.278
New York City	657.49	533.26	0.0641