Supplementary Materials on Clinical Outcomes Database

The clinical outcomes database analyzed in Section 1 aggregates reported data from different hospitals across the world. These hospitals may have different equipment and reporting standards. It was obtained through a human reading process, which is inherently imperfect. For this reason, we now list some important observations and caveats about the database, and refer the reader to our website for a complete list of the 160 papers entered into the database.[1]

- To minimize human error in data reporting, we have verified some key features with additional scrutiny, including mortality, ICU and hospital length of stay, key symptoms (fever, cough, short breath, fatigue, diarrhea) and common comorbidities (hypertension, diabetes).
- Across papers, subcohort divisions may follow different criteria, including: severity of disease (severe vs. mild), mortality (survivors vs. non-survivors), treatment (intubation vs. non-intubation), comorbidity (diabetic vs. non-diabetic). To retain a large enough number of studies in each category, we classify a population as "mild" if the study classifies it as "not asymptomatic" and "mild", "general", or "non-ICU" and not "severe/critical"; and we classify a population as "severe" if the study classifies it as "severe", "critical", "ICU only" or "non-survivors only".
- Studies in this dataset do not always have the same purpose. For instance, many papers from Italy seem to report data only on non-survivors. In addition, some studies focus on the disease's contagion profile, with little information on mortality, discharge, stay length. Data points from these studies may exhibit a high proportion of missing features.
- We have tried to report all lab values in consistent units. We have included a companion document (Reference Ranges) with corresponding reference ranges to facilitate analysis. There are some instances where the reported lab units seem inconsistent with the expected ranges (e.g. for D-Dimer), but we have generally reported the raw values from the source papers.
- The papers entered in the database do not consistently report confidence intervals alongside population means. For this reason, we have declined to provide confidence intervals for the quantities estimated in this part of the paper.
- We intend to continuously update the database as new papers become available. For this reason, the average values reported in this paper may change as more data becomes available.

Table S1: Count and prevalence of treatments among COVID-19 patients, broken down per continent (Asia, Europe, North America). A "-" indicates that fewer than 100 patients in a subpopulation reported on this symptom.

Treatment	Asia		Europe		North America	
	No. Report	Prev. (%)	No. Report.	Prev. (%)	No. Report.	Prev. (%)
Kaletra	5,665	35.2%	–	–	–	–
Oseltamivir	5,901	25.1%	–	–	–	–
Remdesivir	337	47.4%	–	–	868	10.3%
Arbidol	5,902	34.8%	–	–	–	–
Interferon	3,647	51.8%	–	–	–	–
Hydroxychloroquine	6,008	0.7%	–	–	1,235	61.7%
Invasive Ventilation	7,945	8.0%	75,120	4.8%	5,840	19.3%
Proj. Mortality	12,820	16.7%	79,750	9.9%	19,060	15.8%

Table S1 reports statistics on treatments in different continents. The data are reported at a higher level of granularity in early studies in Asia, which hinders direct comparisons. Still, we observe significant differences in the use of hydroxychloroquine and ventilation between Asia, Europe and North America.

[1] https://www.covidanalytics.io/dataset

Supplementary Materials on Clinical Risk Calculators

Method Details

We construct the feature space by aggregating all clinical features for each of the cohorts. We restricted the features to those that have at most 40% of missing values in both datasets (ASST Cremona and HM Hospitals). Missing values are imputed using k-nearest neighbors imputation method (Troyanskaya et al. 2001). The mortality model consists of 22 features. The infection model has a larger feature space, since we are not limited to common features in both datasets. We restrict this model to the 20 most important features, as determined by the algorithm, to ensure usability and reduce the data entry burden on end-users.

We train models for each of the two outcomes of interest (mortality and infection), both including and excluding lab values. This results in a total of four models, referred to as "mortality with lab", "mortality without lab", "infection with lab", and "infection without lab". We use the XGBoost algorithm to train all models (**?**). We leverage a Bayesian optimization framework to select the best model parameters, using the mean cross-validation area under the curve (AUC) across 40 random seeds as the loss function. This technique results in a more accurate tuning compared to standard grid search, yielding better performance on the test set. We use Scikit-learn (Pedregosa et al. 2011) to interface XGBoost and Scikit-optimize (**?**) to perform the hyperparameter tuning. We tune the following parameters for every model: learning rate, γ, λ, α, minimum child weight, maximum tree depth, number of estimators, and the subsample ratio of columns when constructing each tree. All remaining hyperparameters are set to their default value.

Performance Evaluation

Figure S1(a) reports the average receiver operating curve and precision-recall curve for each model. The results are averaged across models generated from 40 random seeds. The mortality models have higher average AUCs than the infection models, although the infection models are stronger when evaluated on precision and recall. As expected, predictive performance deteriorates when lab values are excluded. Yet, the models without lab values still achieve strong performance. In particular, the AUC of mortality model drops only moderately when lab values are excluded. Both models see a similar loss in precision/recall when lab values are excluded.

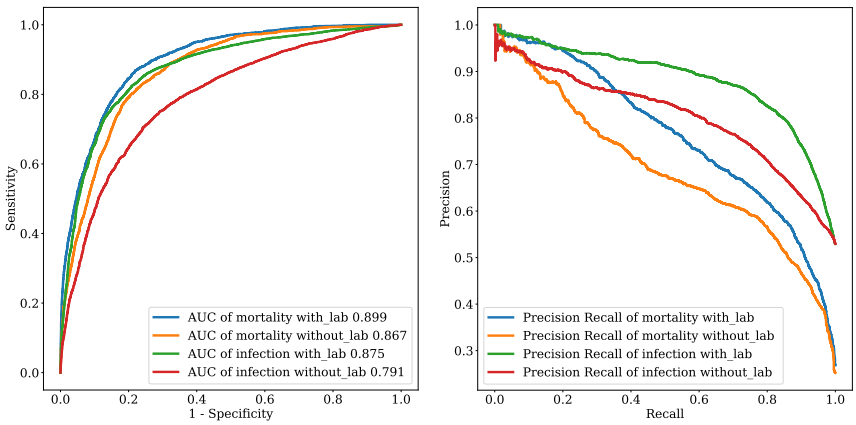

(a) Average ROC Curves

(b) Average Precision Recall Curves

Fig. S1: Bootstrapped results on ROC and Precision Recall curves for all calculators on the testing set.

Table S2 reports threshold-based metrics, which evaluate the discriminative ability of the calculators at a fixed cutoff. We ensure a sensitivity of at least 90% to reflect the high cost of false negatives (missing a death or an infection). We select the highest corresponding threshold to maximize specificity. The results show that the accuracy of the models spans 65%–80%. The mortality calculator with lab values achieves a specificity of 76%. The infection model with lab values has lower specificity (63%), but better precision (74% vs. 56%).

Finally, Figure S2 displays calibration plots, showing the true event rates as a function of the average predicted probabilities. The x-axis bins the population by average predicted risk, and the y-axis plots the true event rate (percentage of deaths or infection). All four risk calculators are well calibrated across subgroups, as the fits are close to the 45-degree line. The bottom plot shows the distribution of predicted risk values from the models. For the mortality calculators, the mean predicted values fall below 10% for most samples, whereas the infection calculators distribute the risk more evenly across the cohort. This reflects the fact that mortality is less prevalent than infection.

Table S2: Performance evaluation summary. Average results across 40 random seeds are reported along with 95% confidence intervals. A minimum threshold of 90% sensitivity is enforced.

Model Type	Lab Values	Threshold	Accuracy	Sensitivity	Specificity	Precision	Negative predictive value	False positive rate
Mortality	Present	17.45 (15.61,19.29)	79.29 (77.47,81.11)	90.38 (90.38,90.38)	75.66 (73.25,78.07)	55.86 (53.46,58.25)	95.97 (95.84,96.1)	24.34 (21.93,26.75)
Mortality	Absent	12.66 (11.11,14.2)	70.85 (68.45,73.24)	90.38 (90.38,90.38)	64.53 (61.37,67.7)	46.15 (43.96,48.33)	95.3 (95.04,95.55)	35.47 (32.3,38.63)
Infection	Present	28.32 (26.63,30.02)	77.58 (76.48,78.68)	90.36 (90.36,90.36)	63.24 (60.9,65.58)	73.59 (72.33,74.85)	85.26 (84.79,85.73)	36.76 (34.42,39.1)
Infection	Absent	27.51 (26.55,28.47)	66.31 (65.37,67.25)	90.36 (90.36,90.36)	39.32 (37.33,41.32)	62.64 (61.85,63.44)	78.12 (77.28,78.96)	60.68 (58.68,62.67)

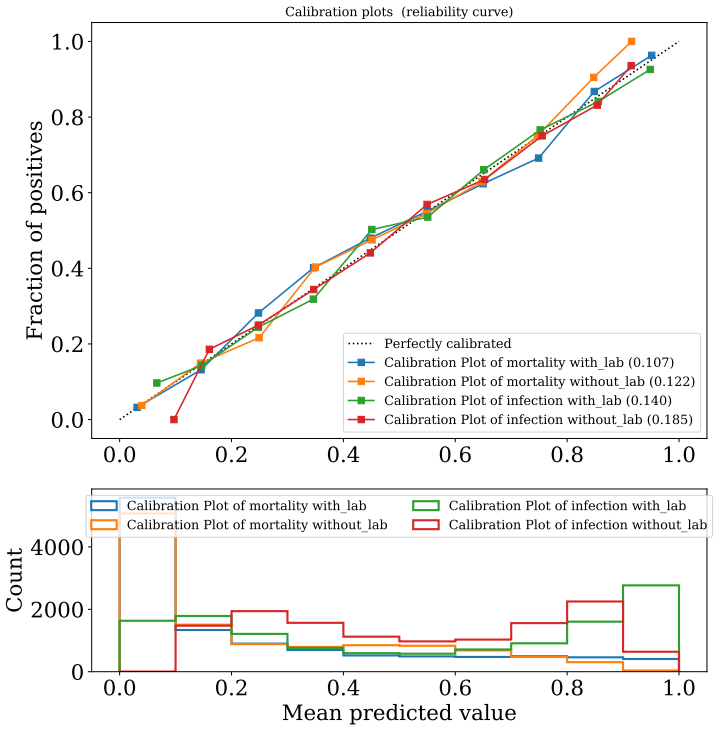

Fig. S2: Bootstrapped Results on the Calibration Curves for both risk calculators on the testing set. The intervals are: [0,10%], (10,20%], (20,30%], …, (90,100%]. The event rates are plotted against the bin mid-points. An ideal event rate is marked by the dotted 45 degree line.

Supplementary Materials on DELPHI-pred and DELPHI-presc

Formulation of DELPHI-pred

General Formulation

The DELPHI model separates people into 11 possible states:

- **Susceptible (S)**: People who have not been infected.
- **Exposed (E)**: People currently infected, but not contagious and within the incubation period.
- **Infected (I)**: People currently infected and contagious.
- **Undetected (U_R) & (U_D)**: People infected and self-quarantined due to the effects of the disease, but not confirmed due to lack of testing. Some of these people recover (U_R) and some die (U_D).
- **Detected, Hospitalized (DH_R) & (DH_D)**: People who are infected, confirmed, and hospitalized. Some of these people recover (DH_R) and some die (DH_D).
- **Detected, Quarantine (DQ_R) & (DQ_D)**: People who are infected, confirmed, and home-quarantined rather than hospitalized. Some of these people recover (DQ_R) and some die (DQ_D).
- **Recovered (R)**: People who have recovered from the disease (and immune).
- **Deceased (D)**: People who have deceased from the disease.

In addition to main functional states, we introduce helper states to calculate a few useful quantities: Total Hospitalized (TH), Total Detected Deceased (DD) and Total Detected Cases (DT). The full mathematical formulation of the model is as follows:

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -\widetilde{\alpha}\gamma(t)S(t)I(t) \tag{1}$$
$$\frac{\mathrm{d}E}{\mathrm{d}t} = \widetilde{\alpha}\gamma(t)S(t)I(t) - \beta E(t) \tag{2}$$
$$\frac{\mathrm{d}I}{\mathrm{d}t} = \beta E(t) - r_d I(t) \tag{3}$$
$$\frac{\mathrm{d}U_R}{\mathrm{d}t} = r_d(1-\widetilde{\mu})(1-p_d)I(t) - \sigma U_R(t) \tag{4}$$
$$\frac{\mathrm{d}DH_R}{\mathrm{d}t} = r_d(1-\widetilde{\mu})p_d p_h I(t) - \kappa DH_R(t) \tag{5}$$
$$\frac{\mathrm{d}DQ_R}{\mathrm{d}t} = r_d(1-\widetilde{\mu})p_d(1-p_h)I(t) - \sigma DQ_R(t) \tag{6}$$
$$\frac{\mathrm{d}U_D}{\mathrm{d}t} = r_d\widetilde{\mu}(1-p_d)I(t) - \widetilde{\tau}U_D(t) \tag{7}$$
$$\frac{\mathrm{d}DH_D}{\mathrm{d}t} = r_d\widetilde{\mu}p_d p_h I(t) - \widetilde{\tau}DH_D(t) \tag{8}$$
$$\frac{\mathrm{d}DQ_D}{\mathrm{d}t} = r_d\widetilde{\mu}p_d(1-p_h)I(t) - \widetilde{\tau}DQ_D(t) \tag{9}$$
$$\frac{\mathrm{d}TH}{\mathrm{d}t} = r_d p_d p_h I(t) \tag{10}$$
$$\frac{\mathrm{d}DD}{\mathrm{d}t} = \widetilde{\tau}(DH_D(t) + DQ_D(t)) \tag{11}$$
$$\frac{\mathrm{d}DT}{\mathrm{d}t} = r_d p_d I(t) \tag{12}$$
$$\frac{\mathrm{d}R}{\mathrm{d}t} = \sigma(U_R(t) + DQ_R(t)) + \kappa DH_R(t) \tag{13}$$
$$\frac{\mathrm{d}D}{\mathrm{d}t} = \widetilde{\tau}(U_D(t) + DQ_D(t) + DH_D(t)) \tag{14}$$

This set of differential equations comprises 11 explicit parameters, defined below. The parameters with a tilde are the parameters that are fitted against historical data for each state; the others are fixed parameters that we estimate using our clinical outcomes database (Section 1).

- $\widetilde{\alpha}$ is the baseline infection rate.
- $\gamma(t)$ measures the effect of government response and is defined as:

$$\gamma(t) = \frac{2}{\pi}\arctan\left(\frac{-(t-\widetilde{t_0})}{\widetilde{k}}\right) + 1,$$

where the parameters $\widetilde{t_0}$ and $\widetilde{k}$ capture, respectively, the timing and the strength of the response. The effective infection rate in the model is $\widetilde{\alpha}\gamma(t)$, which is time dependent.

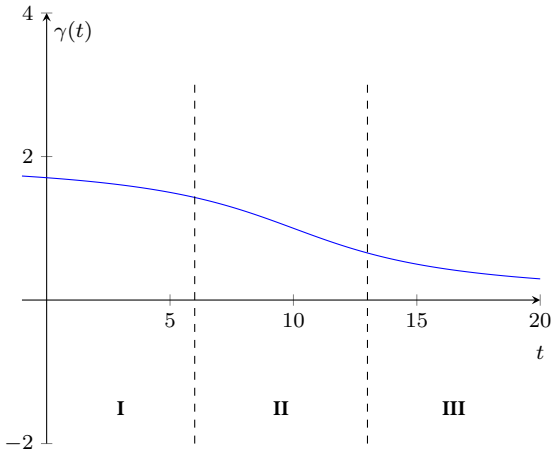

Fig. S3: Illustration of $\gamma(t) = \frac{2}{\pi}\arctan(-\frac{t-10}{5}) + 1$ (i.e., $\widetilde{t_0} = 10$ and $\widetilde{k} = 5$).

- r_d is the rate of detection. This equals to $\frac{\log 2}{T_d}$, where T_d is the median time to detection (fixed to be 2 days). (**?**)
- β is the rate of infection leaving incubation phase. This equals to $\frac{\log 2}{T_\beta}$, where T_β is the median time to leave incubation (fixed at 5 days). (**?**)
- σ is the rate of recovery of non-hospitalized patients. This equals to $\frac{\log 2}{T_\sigma}$, where T_σ is the median time to recovery of non-hospitalized patients (fixed at 10 days). (**??**)
- κ is the rate of recovery under hospitalization. This equals to $\frac{\log 2}{T_\kappa}$, where T_κ is the median time to recovery under hospitalization (fixed at 15 days). (**??**)
- $\widetilde{\tau}$ is the rate of death. This captures the speed at which a dying patient dies, and thus inversely proportional to how long a dying patient stays alive.
- $\widetilde{\mu}$ is the mortality percentage. This is the percentage of people who die from the disease in a particular region. Note this quantity is independent from the rate of death.
- p_d is the (constant) percentage of infectious cases detected. This is set to 20%. (**???**)
- p_h is the (constant) percentage of detected cases hospitalized. This is set to 15%. (**??**)

Therefore, we fit on 5 parameters from the list above ($\widetilde{\alpha}, \widetilde{\mu}, \widetilde{\tau}, \widetilde{t_0}, \widetilde{k}$). In addition, we introduce two additional parameters $\widetilde{k_1}, \widetilde{k_2}$ to account for the unknown initial population in the infected (I) and exposed (E) states (see Supplementary Materials for details). We thus fit seven parameters per area, using the methods described in the core of the paper.

Modeling Government Response

As governments respond to the spread of the epidemic, the rate of infection decreases. We model this by multiplying an initial infection rate with an inverse tangent function, which captures three phases of government response (Figure S3).

- **Phase I**: This phase models the initial response when the government has just started to consider implementing policies. Some people have already changed their behavior in response to early reports, but most people continue business-as-usual activities.
- **Phase II**: This phase is characterized by the sharp decline in infection rate as policies get broadly implemented.
- **Phase III**: This phase reflects the diminishing marginal returns in the decline of the infection rate as the measures reach saturation.

Using parameters $\widetilde{t_0}$ and $\widetilde{k}$, we control the start time and the strength of the measures. We can therefore interpret $\widetilde{t_0}$ as the median day of action, and $\widetilde{k}$ as the median rate of action. This formulation allows us to model, under the same framework, a wide variety of policies—spanning school closures, restriction on mass gatherings, stay-at-home policies, etc.

DELPHI-pred Validation

Figure S4 shows the projected number of deaths in the United States, with projections made on three different weeks, against historical observations. This complements the corresponding figure in the main text reporting the number of projected vs. actual cases. We see that we were generally able to predict the number of deaths up to 4 weeks ahead with good accuracy. One exception is our prediction made on April 3, which is due to a lack of state-level data on deaths at the time (hence, we had to assume a constant mortality rate per state). But after that, our projections have closely followed historical trends.

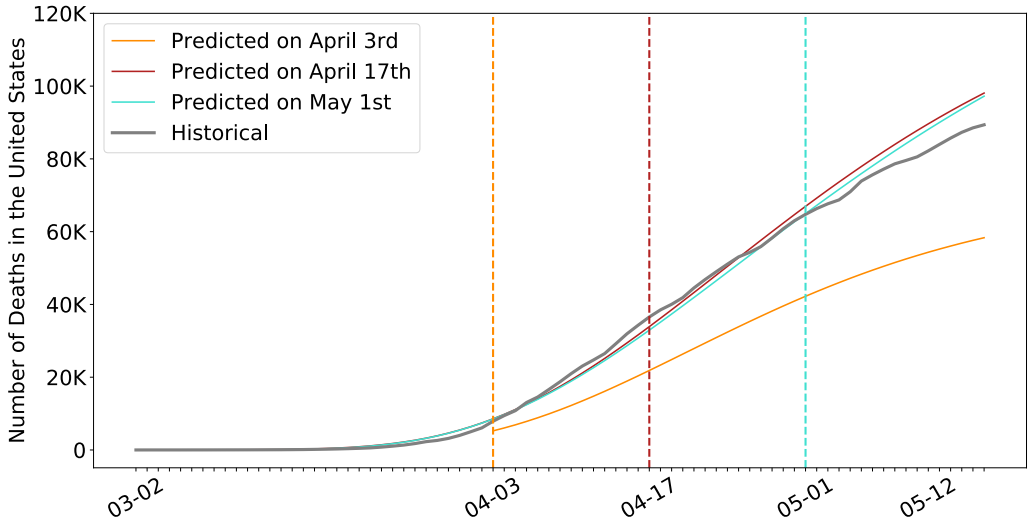

Fig. S4: Cumulative number of deaths in the United States according to our projections made at different points in time, against actual observations

Formulation of DELPHI-pres

Modeling of the Impact of Government Response

Recall that we fit a machine learning model to predict the value of $\gamma(t)$ (fitted by DELPHI-pred), as a function of the policies in place. The objective is to evaluate the impact of each policy on the infection rate in order to simulate its overall effect on the dynamics of the pandemic. Figure S5 shows the resulting regression tree, using state-level data in the United States. The results show that more stringent policies result in lower values of $\gamma(t)$, hence in lower infection rates. For instance, in states with no measure in place, the predicted value of $\gamma(t)$ is 1.304; in states where a stay-at-home policy is in place, the predicted value of $\gamma(t)$ is 0.312; in states where partial social distancing policies are in place, the predicted value of $\gamma(t)$ falls in between.

The main objective of DELPHI-presc is to modify the value of $\gamma(t)$ in DELPHI-pred to account for future changes in social distancing policies, using the values predicted by the tree shown in Figure S5. To this end, we define the following quantities:

- t_c is the time of the policy change.
- k_0 is a *normalized* pair-wise difference between the predicted values of $\gamma(t)$ between policies (with respect to the largest predicted value of $\gamma(t)$ under no measure). For instance, transitioning from stay-at-home to no measure induces an offset $(1.304 - 0.312)/1.304 = +0.761$. All values can be found in Table S3—the offset is positive if the new policy is more lenient, and negative if it is more stringent.
- p_0 is the normalized value of the current policy.

We then correct the government response as follows:

$$\gamma'(t) = \max\left\{\frac{2}{\pi}\arctan\left(-\frac{t-\widetilde{t_0}}{k}\right) + 1 + k_0 \cdot \min\left[\frac{2-\gamma(t_c)}{1-p_0}, \frac{\gamma(t_c)}{p_0}\right], 0\right\}, \quad \forall t \geq t_c.$$

For example, if we are currently in Lockdown and are moving to No measure, then we obtain $k_0 = 0.787$, $p_0 = 0.329/1.544 = 0.213$ and $\gamma(t_c) = \frac{2}{\pi}\arctan\left(-\frac{t_c-\widetilde{t_0}}{k}\right) + 1$.

Application to Policy Assessment

To assess any policy, we run the DELPHI-pred model (governed by the system of differential equations), using the value of the infection rate derived above. We report the impact of the different policies on the case count in the main body of the paper. Figure S6 provides a similar visualization of the effect on the death count in the state of New York. We can draw similar observations regarding the impact of the various policies and the impact of the timing of these policies.

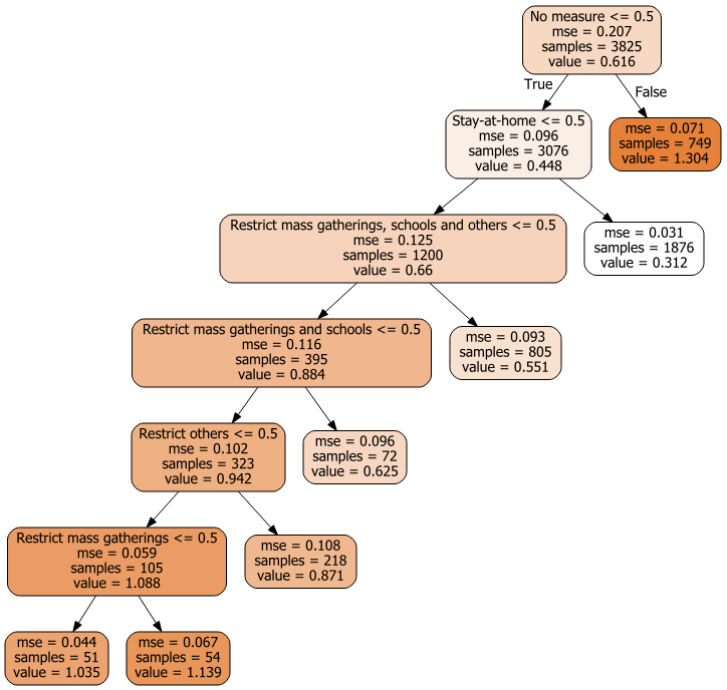

Fig. S5: Regression Tree (CART) predicting an average value of $\gamma(t)$ for each policy.

Table S3: Values of normalized offset computations to correct the estimation of $\gamma(t)$. Policies include: (i) *No measure* ("none"); (ii) *Restrict mass gatherings* ("R-MG"); (iii) *Restrict others* ("R-O"); (iv) *Authorize schools, restrict mass gatherings and others* ("R-MG-O"); (v) *Restrict mass gatherings and schools* ("R-MG-S"); (vi) *Restrict mass gatherings, schools and others* ("R-MG-S-O"); and (vii) *Stay-at-home* ("SAH").

from/to	none	R-MG	R-O	R-MG-O	R-MG-S	R-MG-S-O	SAH
none	0	-0.127	-0.332	-0.206	-0.521	-0.577	-0.761
R-MG	+0.127	0	-0.206	-0.080	-0.294	-0.451	-0.634
R-O	+0.332	+0.206	0	+0.126	-0.189	-0.245	-0.429
R-MG-O	+0.206	+0.080	-0.126	0	-0.314	-0.371	-0.554
R-MG-S	+0.521	+0.294	+0.189	+0.314	0	-0.057	-0.240
R-MG-S-O	+0.577	+0.451	+0.245	+0.371	+0.057	0	-0.183
SAH	+0.761	+0.634	+0.429	+0.554	+0.240	+0.183	0

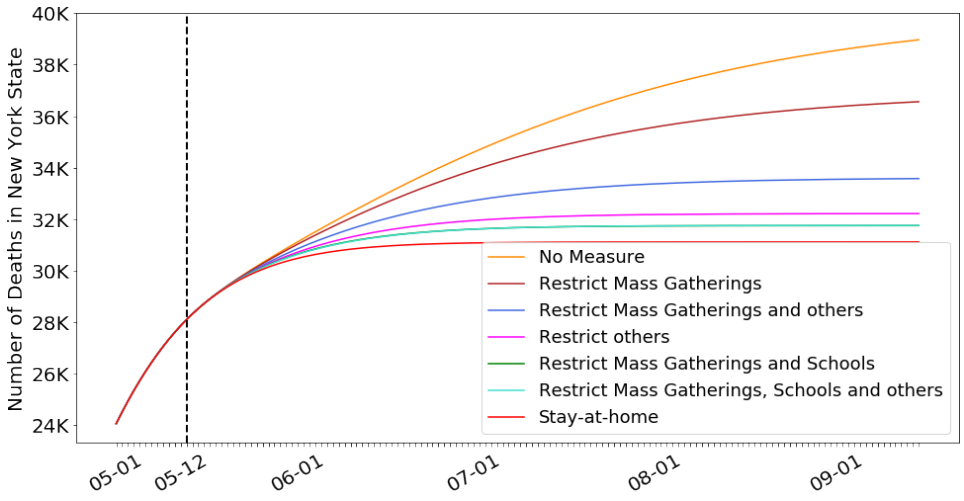

(a) Impact of different policies on the future number of deaths, in NY

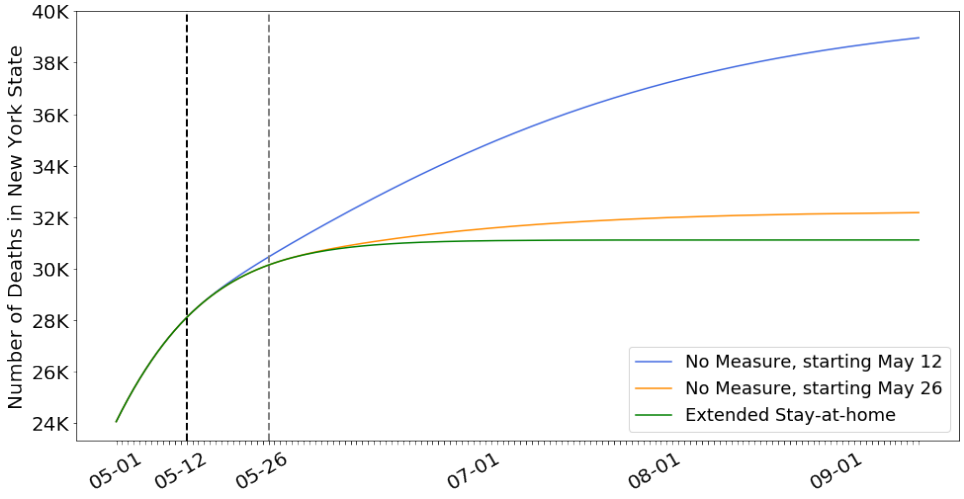

(b) Impact of the timing of policies on the future number of deaths, in NY.

Fig. S6: Impact of different policies on the future number of deaths in the State of New York, for different policies and policy start dates.

Supplementary Materials on Ventilator Allocation

We now detail the formulation of the optimization model proposed for ventilator allocation. We begin by specifying the model mathematically, then discuss data sources and parameter calibration.

Formulation

We consider S states, indexed by $s = 1, \cdots, S$, and D days, indexed by $d = 1, \cdots, D$.

Data.
In formulating the problem, we consider the following data as given:

- $v_{s,d}$ is the demand for ventilators in state s on day d.
- b_s is the base supply of ventilators starting in each state s.
- n_d is the surge supply of ventilators distributed by the federal government on day d.
- $d_{s,s'}$ is the distance between state s and state s'.
- $\tau_{s,s'}$ is the lead time between state s and state s'.

We note two comments regarding these inputs. First, the surge supply n_d corresponds to the number of ventilators that are actually distributed by the federal government on day d: the details of managing the federal stockpile fall beyond the scope of this model. More generally, n_d represents supply available from any exogenous, centralized source. Second, we consider distances between states as a way to encourage transfers between neighboring states. We also calibrate the distances such that $d_{s,s} > 0$ for each state s, to ensure we do not propose meaningless transfers from state s to itself.

Decisions.
We define integer decision variables as follows:

- $x_{s,d} \in \mathbb{Z}^+$ is the supply of ventilators in state s on day d.
- $y_{s,s',d} \in \mathbb{Z}^+$ is the number of ventilators sent from state s to state s' on day d.
- $z_{s,d} \in \mathbb{Z}^+$ is the additional supply state s receives from the federal government on day d.
- $w_{s,d} \in \mathbb{Z}^+$ is the shortage of ventilators in state s on day d relative to the demand $v_{s,d}$.
- $\Delta_{s,d} \in \mathbb{Z}^+$ is the shortage of ventilators in state s on day d relative to the demand with a buffer.

Parameters.
We define the following parameters, which control different policy trade-offs:

- $f_{\max} \in [0, 1]$ is the maximum fraction of its base supply that each state is willing to share. A value of $f_{\max} = 0$ indicates that states are not willing to share any ventilator with other states; a value of $f_{\max} = 1$ indicates that states are willing to share all their ventilator supply with other states.
- $\alpha \in [0, \infty)$ is the percentage of projected demand that states would like to plan for with a supply buffer. For example, $\alpha = 0.1$ will penalize any solution such that supply falls within 10% of projected demand.
- $\lambda \in [0, \infty)$ is a regularization parameter that captures the trade-off between the financial and logistical cost of interstate transfers with the public health cost of ventilator shortages.
- $t_{\min} \in \mathbb{Z}^+$ is the number of days a ventilator is in use after it is shipped to a new location, allowing to control for excessive transfers.
- $\rho \in [0, 1]$ is a relative cost parameter capturing the relative importance of projected shortages vs. worst-case shortages. Each unit of supply that falls short of the projected demand is assigned a cost of 1. Each unit of supply that exceeds the demand but does not exceed the state's desired supply buffer is assigned a cost of ρ.

Objective.
The problem of allocating scarce resources in a pandemic is complex because of the necessity to balance competing interests. We identify two key operational goals: improving public health outcomes, and reducing financial cost. We therefore formulate the ventilator allocation problem with two objectives, minimizing total shortage costs as well as total ventilator transfer costs. Each unit of ventilator shortage is assigned a weight of 1 (for shortage relative to the projected demand) or a weight of $\rho \leq 1$ (for shortage relative to the buffered demand). We formalize the bi-objective problem by means of a penalty on transfers, weighted with a trade-off parameter λ.

$$\min \quad \sum_{s=1}^{S} \sum_{d=1}^{D} (w_{s,d} + \rho \Delta_{s,d}) + \lambda \sum_{s=1}^{S} \sum_{s'=1}^{S} \sum_{d=1}^{D} d_{s,s'} y_{s,s',d}. \tag{15a}$$

Note that ventilators distributed by the federal government are not penalized, as our model simply treats this source as exogenous.

Constraints.

– Initial supply for each state s:

$$x_{s,0} = b_s, \quad \forall s = 1, \cdots, S. \tag{15b}$$

– The supply in each state s on each day d remains higher than the fraction of its initial supply the state wants to retain:

$$x_{s,d} \geq (1 - f_{\max}) b_s, \quad \forall s = 1, \cdots, S,\ d = 1, \cdots, D. \tag{15c}$$

– The transfers from the federal government cannot exceed the available surge supply on each day d:

$$\sum_{s=1}^{S} z_{s,d} \leq n_d, \quad \forall d = 1, \cdots, D. \tag{15d}$$

– The shortage variable corresponds to the positive part of the difference between ventilator demand and supply, if positive, for each state s and day d:

$$w_{s,d} \geq v_{s,d} - x_{s,d}, \quad \forall s = 1, \cdots, S,\ d = 1, \cdots, D. \tag{15e}$$

– The buffer shortage variable is defined such that the total (actual plus buffer) shortage corresponds to the difference between buffered demand and ventilator supply, if positive, for each state s and day d:

$$w_{s,d} + \Delta_{s,d} \geq (1 + \alpha) v_{s,d} - x_{s,d}, \quad \forall s = 1, \cdots, S,\ d = 1, \cdots, D. \tag{15f}$$

– (Conservation of flow) For each state s and day d, today's supply is equal to yesterday's supply plus what is received today from the government and the other states, minus what is sent to other states, with $\tau_{s',s}$ reflecting shipments' lead times:

$$x_{s,d} = x_{s,d-1} + z_{s,d} + \sum_{s'=1}^{S} y_{s',s,d-\tau_{s',s}} - \sum_{s'=1}^{S} y_{s,s',d}, \quad \forall s = 1, \cdots, S,\ d = 1, \cdots, D. \tag{15g}$$

– (Minimum days in use) For each state s and day d, any incoming ventilator, either from another state or from the federal government, cannot be shipped out for at least t_{min} days. This constraint ensures that ventilators are not transferred too often.

$$\sum_{d'=\max(1,d-t_{\min})}^{d-1} \left(z_{s,d'} + \sum_{s'=1}^{S} y_{s',s,d'} \right) \leq x_{s,d}, \quad \forall s = 1, \cdots, S,\ d = 1, \cdots, D. \tag{15h}$$

– Any state s facing a shortage on day d cannot ship any ventilators to other states on day d. To write this constraint, we define auxiliary binary variables $a_{s,d} \in \{0, 1\}$ indicating if there is a shortage in state s on day d and a parameter $V_{\max}$ providing a trivial upper bound on the number of ventilators a state can ship per day (we use a value of 3,000 which does not restrict the solution, as it exceeds the shortage faced by any state on any given day).

$$w_{s,d} + \Delta_{s,d} \leq v_{s,d}(1 + \alpha) a_{s,d}, \tag{15i}$$

$$\sum_{s'=1}^{S} y_{s,s',d} \leq V_{\max}(1 - a_{s,d}). \tag{15j}$$

Data Sources and Parameter Calibration

Our optimization model is complex enough to model high-level dynamics of scarce resource allocation, yet simple enough to only require simple data inputs. We now describe our methodology in collecting the key data necessary to solve this optimization problem.

Demand.

The most important input data is the forecasted ventilator demand $v_{s,d}$. Consistent with our end-to-end data-driven approach, and in contrast with other ventilator allocation approaches (Mehrotra et al. 2020), we develop our own demand forecasts using DELPHI-pred (Section 3 of the main text). Recall that DELPHI-pred does not only estimate the number of cases, but also the number of hospitalizations, equal to $DH_R + DH_D$. We then apply a 25% ratio to estimate the number of ventilators in use—given that, in our clinical outcomes database, 25% of hospitalized patients are on a ventilator. Ultimately, we can use the DELPHI-pred outputs to derive projections of ventilator demand at the state level and at the daily level—consistently with the optimization input $v_{s,d}$.

Supply.
It can be difficult to estimate how many ventilators are available in each state as well as at the federal level. For the base supply b_s, we use inventory levels from a 2010 American Medical Association report (**?**). We adjust this number for population growth, under the assumption that the number of ventilators per capita has remained constant in each state.

Of course, ventilators can also be used to treat non-COVID-19 patients. We assume that 50% of ventilator supply across the board is unavailable due to non-COVID-19 usage, in line with other estimates (Mehrotra et al. 2020).

In addition, our model takes into account the daily availability n_d of ventilators at the federal level. Estimating this quantity from publicly-available sources is both difficult due to limited data and politically fraught. We use the estimate from the Society for Critical Medicine that the federal stockpile contains at least 12,700 ventilators (**?**). Some news reports suggest a lower estimate of 10,000 based on some defects in the stockpiled equipment (**?**), while others suggest an estimate of 16,600 based on older model repairs (**?**). Based on these reference points, we estimate roughly 13,500 available ventilators and assume that they can be deployed evenly over a month. In other words, we allow 450 ventilators to be deployed each day for 30 days (starting on day 4 to allow for lead times). This gradual release reflects potential operational constraints and strategic considerations of controlling the release of inventory in case of unexpected outbreaks. Yet, given the underlying uncertainty, we perform sensitivity analysis to explore how the model's recommendation varies with the federal stockpile.

Distances and lead times.
We compute the interstate distance $d_{s,s'}$ as the Euclidean distance between the centers of states s and s', and we let the lead time parameter $\tau_{s,s'}$ equal 3 days for every pair of states. Our choice of a conservative uniform lead time for shipments is motivated by simplicity concerns. This could be improved, in future work, to better reflect efficiencies in the US shipping infrastructure.

Trade-off parameters.
- Understanding the impact of states' willingness to share ventilators with other states is a key takeaway from our model. In Figure 7B, we vary the fraction $f_{\max}$ of each state's pooling supply between $\{0\%, 5\%, 10\%, 15\%, 20\%\}$. Results indicate there is little additional efficiency to be gained from states sharing more than 20% of their supply.
- Understanding the relative importance of federal surge supply compared to interstate transfers is another interesting takeaway from our model. In Figure 7A, we show the effects of removing federal surge supply, or preventing interstate transfers, on ventilator shortage reduction. We show more detailed sensitivity analysis results in Section 5.
- We vary the parameter λ to derive the the Pareto-optimal frontier of the trade-off between inter-state transfers vs. ventilator shortages (for instance in Fig. 7B). As λ tends to zero, transfers incur no cost other than rendering the ventilators unavailable during shipment; as λ tends to infinity, transfers become heavily discouraged.
- The parameter α models uncertainty in the demand forecast as well as robustness to inefficiency in ventilator allocation within each state. We vary the percentage α of buffered demand within $\{0\%, 5\%, 10\%, 20\%\}$ in Section 5.

Finally, we choose the following values for the remaining two parameters, which have a small impact on the final solution.

- We set the value of $t_{\min}$ to 10 days, based on the clinical outcomes database (Section 1).
- We set ρ, the cost of shortage with respect to buffered demand relative to the cost of shortage with respect to real demand, to 0.25.

References

Adelman D (2020) Thousands Of Lives Could Be Saved In The US During The COVID-19 Pandemic If States Exchanged Ventilators: Study examines how lives could be saved by allowing US states to exchange ventilators during the COVID-19 pandemic. Health Affairs pp 10–1377

Anderson RM, Heesterbeek H, Klinkenberg D, Hollingsworth TD (2020) How will country-based mitigation measures influence the course of the COVID-19 epidemic? The Lancet 395(10228):931–934

Bein T, Grasso S, Moerer O, Quintel M, Guerin C, Deja M, Brondani A, Mehta S (2016) The standard of care of patients with ARDS: ventilatory settings and rescue therapies for refractory hypoxemia. Intensive care medicine 42(5):699–711

Billingham S, Widrick R, Edwards NJ, Klaus S (2020) COVID-19 (SARS-CoV-2) Ventilator Resource Management Using a Network Optimization Model and Predictive System Demand. medRxiv

Breiman L, Friedman J, Stone CJ, Olshen RA (1984) Classification and regression trees. CRC press

Byrd RH, Gilbert JC, Nocedal J (2000) A trust region method based on interior point techniques for nonlinear programming. Mathematical Programming 89(1):149–185

Caruso D, Zerunian M, Polici M, Pucciarelli F, Polidori T, Rucci C, Guido G, Bracci B, de Dominicis C, Laghi A (2020) Chest CT features of COVID-19 in Rome, Italy. Radiology p 201237

Chen T, Guestrin C (2016) Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, pp 785–794

Cheng Y, Luo R, Wang K, Zhang M, Wang Z, Dong L, Li J, Yao Y, Ge S, Xu G (2020) Kidney disease is associated with in-hospital death of patients with covid-19. Kidney international

Cornejo RA, Díaz JC, Tobar EA, Bruhn AR, Ramos CA, González RA, Repetto CA, Romero CM, Gálvez LR, Llanos O, et al. (2013) Effects of prone positioning on lung protection in patients with acute respiratory distress syndrome. American journal of respiratory and critical care medicine 188(4):440–448

Fernandes N (2020) Economic effects of coronavirus outbreak (COVID-19) on the world economy. Available at SSRN 3557504

Forster P, Forster L, Renfrew C, Forster M (2020) Phylogenetic network analysis of SARS-CoV-2 genomes. Proceedings of the National Academy of Sciences 117(17):9241–9243

Garg S (2020) Hospitalization rates and characteristics of patients hospitalized with laboratory-confirmed coronavirus disease 2019—COVID-NET, 14 states, March 1–30, 2020. MMWR Morbidity and Mortality Weekly Report 69

Gershengorn HB, Hu Y, Chen JT, Hsieh SJ, Dong J, Gong MN, Chan CW (2020) The impact of high-flow nasal cannula use on patient mortality and the availability of mechanical ventilators in covid-19. Annals of the American Thoracic Society (ja)

Goyal P, Choi JJ, Pinheiro LC, Schenck EJ, Chen R, Jabri A, Satlin MJ, Campion Jr TR, Nahid M, Ringel JB, et al. (2020) Clinical characteristics of Covid-19 in New York City. New England Journal of Medicine

Guan Wj, Ni Zy, Hu Y, Liang Wh, Ou Cq, He Jx, Liu L, Shan H, Lei Cl, Hui DS, et al. (2020) Clinical characteristics of coronavirus disease 2019 in china. New England Journal of Medicine 382(18):1708–1720

Holland LA, Kaelin EA, Maqsood R, Estifanos B, Wu LI, Varsani A, Halden RU, Hogue BG, Scotch M, Lim ES (2020) An 81 nucleotide deletion in SARS-CoV-2 ORF7a identified from sentinel surveillance in Arizona (Jan-Mar 2020). Journal of Virology

Huang C, Wang Y, Li X, Ren L, Zhao J, Hu Y, Zhang L, Fan G, Xu J, Gu X, et al. (2020) Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. The Lancet 395(10223):497–506

Huang HC, Araz OM, Morton DP, Johnson GP, Damien P, Clements B, Meyers LA (2017) Stockpiling ventilators for influenza pandemics. Emerging infectious diseases 23(6):914

Johns Hopkins University (2020) Center for Systems Science and Engineering. https://systems.jhu.edu/research/public-health/ncov/

Jordan RE, Adab P, Cheng K (2020) COVID-19: risk factors for severe disease and death

Kampf G, Todt D, Pfaender S, Steinmann E (2020) Persistence of coronaviruses on inanimate surfaces and its inactivation with biocidal agents. Journal of Hospital Infection

Kissler SM, Tedijanto C, Goldstein E, Grad YH, Lipsitch M (2020) Projecting the transmission dynamics of sars-cov-2 through the postpandemic period. Science

Knight SR, Ho A, Pius R, Buchan I, Carson G, Drake TM, Dunning J, Fairfield CJ, Gamble C, Green CA, et al. (2020) Risk stratification of patients admitted to hospital with covid-19 using the ISARIC WHO Clinical Characterisation Protocol: development and validation of the 4C Mortality Score. bmj 370

Lagarias JC, Reeds JA, Wright MH, Wright PE (1998) Convergence properties of the Nelder–Mead simplex method in low dimensions. SIAM Journal on Optimization 9(1):112–147

Lan L, Xu D, Ye G, Xia C, Wang S, Li Y, Xu H (2020) Positive RT-PCR test results in patients recovered from COVID-19. Jama 323(15):1502–1503

LANL (2020) LANL COVID-19 Cases and Deaths Forecasts

Ling W (2020) C-reactive protein levels in the early stage of COVID-19. Medecine et Maladies Infectieuses

Liu Y, Du X, Chen J, Jin Y, Peng L, Wang HH, Luo M, Chen L, Zhao Y (2020) Neutrophil-to-lymphocyte ratio as an independent risk factor for mortality in hospitalized patients with COVID-19. Journal of Infection

Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) Advances in Neural Information Processing Systems 30, Curran Associates, Inc., pp 4765–4774

Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, Katz R, Himmelfarb J, Bansal N, Lee SI (2020) From local explanations to global understanding with explainable AI for trees. Nature Machine Intelligence 2(1):2522–5839

McKibbin WJ, Fernando R (2020) The global macroeconomic impacts of COVID-19: Seven scenarios. CAMA Working Paper

Mehrotra S, Rahimian H, Barah M, Luo F, Schantz K (2020) A model of supply-chain decisions for resource sharing with an application to ventilator allocation to combat COVID-19. Naval Research Logistics (NRL)

Meng L, Qiu H, Wan L, Ai Y, Xue Z, Guo Q, Deshpande R, Zhang L, Meng J, Tong C, et al. (2020) Intubation and Ventilation amid the COVID-19 Outbreak Wuhan's Experience. Anesthesiology: The Journal of the American Society of Anesthesiologists

National Health Service (2020) Check if you have Coronavirus symptoms. https://www.nhs.uk/conditions/coronavirus-covid-19/check-if-you-have-coronavirus-symptoms, (Accessed 11 May 2020)

New York Times (2020) New York Times, Coronavirus in the U.S.: Latest Map and Case Count. https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html

Pathak PA, Sönmez T, Unver MU, Yenmez MB (2020) Leaving no ethical value behind: Triage protocol design for pandemic rationing. Working Paper 26951, National Bureau of Economic Research, DOI 10.3386/w26951, URL http://www.nber.org/papers/w26951

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, et al. (2011) Scikit-learn: Machine learning in Python. The Journal of Machine Learning Research 12:2825–2830

Pei S, Shaman J (2020) Initial Simulation of SARS-CoV2 Spread and Intervention Effects in the Continental US. medRxiv

Petrilli CM, Jones SA, Yang J, Rajagopalan H, O'Donnell LF, Chernyak Y, Tobin K, Cerfolio RJ, Francois F, Horwitz LI (2020) Factors associated with hospitalization and critical illness among 4,103 patients with COVID-19 disease in New York City. medRxiv

Ranney ML, Griffeth V, Jha AK (2020) Critical supply shortages—the need for ventilators and personal protective equipment during the COVID-19 pandemic. New England Journal of Medicine

Ray EL, Wattanachit N, Niemi J, Kanji AH, House K, Cramer EY, Bracher J, Zheng A, Yamana TK, Xiong X, et al. (2020) Ensemble forecasts of coronavirus disease 2019 (covid-19) in the us. MedRXiv

Sanchel S, Ting Lin Y, Xu C, Romero-Severson E, Hengartner N, Ke R (2020) High Contagiousness and Rapid Spread of Severe Acute Respiratory Syndrome Coronavirus 2. Emerging Infectious Diseases 27(7)

Shi Q, Sun J, Zhang W, Zou L, Liu Y, Li J, Kan X, Dai L, Yuan S, Yu W, et al. (2020a) Serum calcium as a biomarker of clinical severity and prognosis in patients with coronavirus disease 2019: a retrospective cross-sectional study. Critical Care and Emergency Medicine

Shi Y, Wang Y, Shao C, Huang J, Gan J, Huang X, Bucci E, Piacentini M, Ippolito G, Melino G (2020b) COVID-19 infection: the perspectives on immune responses

Troyanskaya O, Cantor M, Sherlock G, Brown P, Hastie T, Tibshirani R, Botstein D, Altman RB (2001) Missing value estimation methods for DNA microarrays. Bioinformatics 17(6):520–525, DOI 10.1093/bioinformatics/17.6.520, URL http://dx.doi.org/10.1093/bioinformatics/17.6.520, /oup/backfile/contentpublic/journal/bioinformatics/17/6/10.1093/bioinformatics/17.6.520/2/170520.pdf

US Center for Disease Control (2020) US Center for Disease Control, Symptoms of Coronavirus. https://www.cdc.gov/coronavirus/2019-ncov/symptoms-testing/symptoms.html, (Accessed 11 May 2020)

Vahidy FS, Nicolas JC, Meeks JR, Khan O, Jones SL, Masud F, Sostman HD, Phillips RA, Andrieni JD, Kash BA, et al. (2020) Racial and ethnic disparities in SARS-CoV-2 pandemic: Analysis of a COVID-19 observational registry for a diverse US metropolitan population. medRxiv

Velavan TP, Meyer CG (2020) The COVID-19 epidemic. Tropical medicine & international health 25(3):278

Wang Y, Wang Y, Chen Y, Qin Q (2020) Unique epidemiological and clinical features of the emerging 2019 novel coronavirus pneumonia (COVID-19) implicate special control measures. Journal of medical virology 92(6):568–576

Wilkerson RG, Adler JD, Shah NG, Brown R (2020) Silent hypoxia: A harbinger of clinical deterioration in patients with covid-19. The American Journal of Emergency Medicine

World Health Organization (2020) Coronavirus. https://www.who.int/health-topics/coronavirus, (Accessed 11 May 2020)

Xie J, Covassin N, Fan Z, Singh P, Gao W, Li G, Kara T, Somers VK (2020) Association Between Hypoxemia and Mortality in Patients With COVID-19. Mayo Clinic Proceedings DOI 10.1016/j.mayocp.2020.04.006, URL https://doi.org/10.1016/j.mayocp.2020.04.006.https://linkinghub.elsevier.com/retrieve/pii/S0025619620303670

Yan L, Zhang HT, Goncalves J, Xiao Y, Wang M, Guo Y, Sun C, Tang X, Jing L, Zhang M, et al. (2020) An interpretable mortality prediction model for covid-19 patients. Nature Machine Intelligence pp 1–6

Zhou F, Yu T, Du R, Fan G, Liu Y, Liu Z, Xiang J, Wang Y, Song B, Gu X, Guan L, Wei Y, Li H, Wu X, Xu J, Tu S, Zhang Y, Chen H, Cao B (2020) Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study. The Lancet 395, DOI 10.1016/S0140-6736(20)30566-3, URL https://doi.org/10.1016/