

# Stockpiling Ventilators for Influenza Pandemics

Hsin-Chan Huang,<sup>1</sup> Ozgur M. Araz, David P. Morton, Gregory P. Johnson, Paul Damien, Bruce Clements, Lauren Ancel Meyers

In preparing for influenza pandemics, public health agencies stockpile critical medical resources. Determining appropriate quantities and locations for such resources can be challenging, given the considerable uncertainty in the timing and severity of future pandemics. We introduce a method for optimizing stockpiles of mechanical ventilators, which are critical for treating hospitalized influenza patients in respiratory failure. As a case study, we consider the US state of Texas during mild, moderate, and severe pandemics. Optimal allocations prioritize local over central storage, even though the latter can be deployed adaptively, on the basis of real-time needs. This prioritization stems from high geographic correlations and the slightly lower treatment success assumed for centrally stockpiled ventilators. We developed our model and analysis in collaboration with academic researchers and a state public health agency and incorporated it into a Web-based decision-support tool for pandemic preparedness and response.

Diligent preparation and effective countermeasures are critical to mitigating future influenza pandemics. The 1918 influenza pandemic, the most severe in recent history, resulted in  $\approx 50$  million deaths globally, of which nearly 675,000 occurred in the United States (1). The 1957 and 2009 pandemics were less severe, causing  $\approx 70,000$  and 9,000–18,000 US deaths, respectively (1). The US Department of Health and Human Services (HHS) estimated (2) that 865,000 US residents would be hospitalized during a moderate pandemic (as in 1957 and 1968) and 9.9 million during a severe pandemic (as in 1918).

When severe influenza outbreaks cause high rates of hospitalization, a surge of medical resources is required, including critical care supplies, antiviral medications, and

personal protection equipment. Given uncertainty in the timing and severity of the next pandemic, as well as the time required to manufacture medical countermeasures, stockpiling is central to influenza preparedness (3). However, difficulty in forecasting and limited public health budgets often constrain decisions about sizes, locations, and deployment of such stockpiles.

Mechanical ventilators are essential for treating influenza patients in severe acute respiratory failure. Substantial concern exists that intensive care units (ICUs) might have insufficient resources to treat all persons requiring ventilator support. Prior studies argue that current capacities are insufficient to handle even moderately severe pandemics and that sentinel reporting and model-based decision-making are critical for managing limited resources (4–6). For this reason, the United States has stockpiled mechanical ventilators in strategically located warehouses for use in public health emergencies, such as an influenza pandemic. The Centers for Disease Control and Prevention (CDC) manages this Strategic National Stockpile (SNS) and has plans for rapid deployment to states during critical events (7).

However, SNS ventilators might not suffice to meet demand during a severe public health emergency. In 2002, the SNS included  $\approx 4,400$  ventilators (8,9), and 4,500 SNS ventilators were added during 2009 and 2010. The American Association for Respiratory Care suggested the SNS inventory should increase to at least 11,000–16,000 ventilators in preparation for a severe influenza pandemic (10). The American Association for Respiratory Care and CDC (11) provide training on 3 types of SNS ventilators—LP10 (Covidien, Boulder, CO, USA); LTV1200 (CareFusion, Yorba Linda, CA, USA); and Uni-vent Eagle 754 (Impact Instrumentation, Inc., West Caldwell, NJ, USA)—to ensure proper use nationwide. In addition to the nationally held SNS, some US states maintain their own stockpiles.

Successful deployment of central ventilator stockpiles, whether federal or state, requires rapid distribution to healthcare facilities with patients in need, along with adequate bed space, requisite supplies, and trained personnel

---

Author affiliations: The University of Texas at Austin, Austin, Texas, USA (H.-C. Huang, G.P. Johnson, P. Damien, L.A. Meyers); University of Nebraska, Lincoln, Nebraska, USA (O.M. Araz); University of Nebraska Medical Center, Omaha, Nebraska, USA (O.M. Araz); Northwestern University, Evanston, Illinois, USA (D.P. Morton); Department of State Health Services, Austin (B. Clements); Santa Fe Institute, Santa Fe, New Mexico, USA (L.A. Meyers)

DOI: <http://dx.doi.org/10.3201/eid2306.161417>

<sup>1</sup>Current affiliation: Precima, LoyaltyOne US, Inc., Chicago, IL, USA.

(12–14). Robust methods for sizing and locating ventilator stockpiles have not yet been developed (15). Wilgis (16) discussed the relative merits of central stockpiling of ventilators to be distributed during an emergency versus distributing ventilators to hospitals a priori. Centralized stockpiles benefit from better inventory tracking, more timely repairs, and superior allocation of a limited resource, but hospital-based supplies facilitate staff training, enable immediate use, and avoid the cost and logistical challenges of central storage and deployment.

We developed an optimization framework for allocating mechanical ventilators to central and local stockpiles to ensure adequate surge capacity during a future pandemic. This data-driven method considers the trade-off between risk and stockpiling cost, where risk is measured 2 ways: expected value of unmet demand (EUD; number of influenza patients not receiving required ventilation) and probability of unmet demand (PUD; probability at least 1 patient does not receive required ventilation). For a given set of healthcare providers in a region, we determined the optimal number of mechanical ventilators to stockpile centrally and at each provider site.

As a case study, we considered the US state of Texas under mild, moderate, and severe influenza pandemic scenarios. Based on the Texas Department of State Health Services (DSHS) response to the 2009 influenza A(H1N1) pandemic and planning efforts for future pandemics, we considered stockpiling across 9 sites: a centrally held state stockpile and local stockpiles in each of Texas' 8 health service regions (HSRs; online Technical Appendix Figure 1, <https://wwwnc.cdc.gov/EID/article/23/6/16-1417-Techapp1.pdf>). We implemented this model in a Web-based decision-support tool for DSHS (17).


## Methods

Our approach had 3 stages (Figure 1). First, we estimated the weekly influenza-related hospitalizations at each site using an adaptive time-dynamic forecasting model. Second, we estimated the number of patients requiring ventilation at each site during the peak week on the basis of published estimates of the proportion of hospitalized influenza patients requiring mechanical ventilation. Finally, we allocated ventilators at minimum cost to achieve a specified level of preparedness through a mathematical optimization model. That model assumed centrally stockpiled ventilators have slightly lower treatment rates than locally held ventilators. In the Texas case study, we estimated hospitalizations under a mild scenario by fitting the forecasting model to data from the 2009 influenza A(H1N1) pandemic, and then we scaled the estimates to simulate moderate and severe pandemics. We summarize our optimization model and forecasting methods and provide details in the online Technical Appendix.

### Optimization Model for Ventilator Stockpiling

Using a 2-stage model, we optimized the allocation of ventilators to a central stockpile and several local stockpiles (at healthcare facilities) to ensure that all sites had sufficient surge capacity to manage the peak of an influenza pandemic. We considered the trade-off between unmet ventilator demand (risk) and the cost of stockpiling ventilators (assuming cost is proportional to number of ventilators) and minimized cost while limiting risk to a specified threshold. We analyzed the risk–cost trade-off by solving a family of optimization models, across a range of risk thresholds.

We assumed the following: each stockpiled ventilator is both child- and adult-capable, will be used to treat at most 1 patient during peak demand, and will not be used for



**Figure 1.** Overview of methods for projecting the need to stockpile ventilators for an influenza pandemic, Texas, USA. First, a forecasting model was used to estimate weekly hospitalizations at each site on the basis of historical ILI hospitalization data and CDC ILINet reports. Second, 3 additional factors, along with a spatial correlation coefficient, were used to form a probability distribution for peak-week ventilator demand at each site. Third, an optimization model was solved to determine local and central stockpile allocations and generate trade-off curves between the expected unmet demand and total stockpile and between the probability of unmet demand and total stockpile. CDC, Centers for Disease Control and Prevention; HSR, health service region; ICU, intensive care unit; ILI, influenza-like illness.

noninfluenza patients; stockpiles were established before the pandemic, and centrally held ventilators can be deployed only once to a site with excess demand (i.e., no redeployment is allowed, even though influenza peaks might be asynchronous across sites); patients requiring ventilatory support cannot move between sites; locally held ventilators are immediately and successfully administered to on-site patients requiring care, and centrally held ventilators incur wastage (i.e., a reduced fraction are successfully deployed to healthcare sites upon demand); patients at all sites have equal priority; and consumable ventilator supplies, requisite staffing, and space are in sufficient supply. The optimization model considers expected unmet demand, and we calculated the probability of unmet demand post hoc, as a secondary risk measure.

### Texas Case Study

We assumed that ventilators can be stockpiled centrally by the Texas DSHS or locally by hospitals in Texas' 8 HSRs (online Technical Appendix Figure 1). We further assumed that local stockpiles within an HSR are available throughout the HSR by movement of either ventilators or patients among healthcare facilities; that is, any patient within an HSR requiring ventilatory support has access to available ventilators within that HSR. To model peak ventilator demand across Texas' 8 HSRs under different pandemic scenarios, we 1) estimated the region-to-region (HSR-to-HSR) correlation in peak-week ventilator demand on the basis of 2003–2008 seasonal influenza hospitalization data and 2009 pandemic hospitalization data; 2) generated probabilistic estimates of peak-week influenza-related hospitalizations by fitting our forecasting model to a baseline (mild) pandemic scenario estimated from 2009 pandemic data; 3) used the estimates derived in steps 1 and 2 to estimate the numbers of influenza patients requiring mechanical ventilation at the pandemic peak in each HSR; and 4) generated moderate and severe pandemic scenarios by scaling the peak demand estimates of the mild scenario. We summarize the parameters we used to estimate peak ventilator demand under different pandemic scenarios (Table 1) and outline the data and methods used to estimate these parameters.

### Texas Influenza Data

We obtained weekly Texas hospital discharge data for 2003–2009, filtered for International Classification of Diseases,

Ninth Revision, codes 487 and 488, corresponding to influenza-like-illness (ILI), and aggregated by HSR. These data comprised all Texas hospitals except those in counties with populations <35,000, those with <100 hospital beds, and those that do not accept insurance or government reimbursement. The number of ILI-related hospital discharges during the 2009 pandemic (April–December 2009) totaled 29,459. We assessed the validity of this International Classification of Diseases, Ninth Revision–based filter for influenza through comparison with CDC (18) and Texas DSHS (19) reports.

We also analyzed data from the CDC ILINet, which tracks weekly outpatient visits related to ILI. CDC guidelines define ILI as fever of at least 100°F and cough and/or sore throat in the absence of a known cause other than influenza. A network of 2,400 sites (health departments, laboratories, vital statistics offices, healthcare providers, and emergency departments) in the 50 states reports to ILINet, and we obtained weekly reports during the 2009 H1N1 pandemic for Texas, aggregated by HSR. Finally, Texas DSHS provided data on the 3,730 ventilators stockpiled in Texas in 2009 (online Technical Appendix Table 1).

### Region-to-Region Correlation in Peak Hospitalizations

For each of the 6 influenza seasons in years starting 2003–2008 and the 2009 pandemic, we calculated peak-week ILI hospitalizations requiring ventilation in each HSR. Across all 28 pairs of HSRs, the average correlation in peak ventilator demand was  $0.72 \pm \text{SD } 0.23$  (range 0.22–0.98). One HSR, with <3% of total hospitalizations during 2009, had pairwise correlations as low as 0.22, but all other pairs of HSRs had coefficients >0.50. We found similar spatiotemporal correlations in hospitalizations when we estimated pairwise HSR-to-HSR correlations for various values of the proportion of ventilated patients requiring 2 weeks (rather than 1 week) of ventilation, and weekly numbers of ILI hospitalizations requiring ventilation, throughout the 2003–2008 influenza seasons and the 2009 pandemic. Given this consistent statewide synchrony in epidemic intensity, we made the simplifying assumption that peak hospitalizations in all HSRs were correlated at a pairwise level of 0.70.

### Forecasting Model for Hospitalizations

We used a dynamic linear forecasting model (online Technical Appendix), which provides a powerful method for

**Table 1.** Parameters for estimating peak-week ventilator demand in mild, moderate, and severe influenza pandemics, Texas, USA\*

Parameter	Mild (2009-like)	Moderate (1957- and 1968-like)	Severe (1918-like)	Source
Hospitalization scaling over mild	1	3.14	36	(2,21)
Proportion of hospitalized ILI patients requiring ICU care	0.2	0.25	0.25	(2,19,21,22,23)
Proportion of ICU patients requiring ventilation	0.5	0.5	0.5	(2,21,22)
Proportion of ventilated patients requiring 2 weeks of ventilation	0.4	0.4	0.4	(22)
Region-to-region correlation for peak-week demand	0.7	0.7	0.7	Estimated

\*ICU, intensive care unit; ILI, influenza-like illness.

capturing system uncertainty when numerous dynamic factors influence a system (20). Although hospitalizations could be forecast only on the basis of historical ILI data, our approach can incorporate additional predictors, such as the most recent ILINet reports, to better represent demand uncertainty. Our forecasting method estimated weekly influenza-related hospitalizations in the 8 HSRs for 2009 pandemic-like scenarios, using CDC ILINet influenza A(H1N1)pdm09 weekly reports as a predictor, from the week ending April 4, 2009, through the week ending December 26, 2009. To account for seasonality, we assumed 5 distinct time periods (September–October, November–December, January–February, March–April, and May–August). We also considered other candidate variables, such as school calendars, humidity, and Google Flu Trends, but these did not substantially improve peak estimates.

### Estimating Regional Ventilator Demand

To estimate regional ventilator peak-week demand, we integrated our weekly forecasts of influenza hospitalizations in each region, the spatial correlation in peak-week demand for ventilators, and 3 additional factors: 1) the proportion of hospitalized ILI patients requiring ICU care, 2) the proportion of ICU patients requiring ventilation, and 3) the proportion of ventilated patients requiring 2 weeks of ventilation (rather than 1). To model “spillover” demand of patients requiring 2 weeks of ventilation, we used week-to-week correlations in influenza hospitalization (online Technical Appendix Table 2).

### Proportion of Hospitalized ILI Patients Requiring ICU Care

From 2009 influenza hospital discharge data, we estimated that 18% of patients required ICU care during the peak week. Texas DSHS reported that 23% of the 2,030 confirmed influenza A(H1N1)pdm09 patients requiring hospitalization in Texas during October–December 2009 required ICU care (19). For moderate and severe planning scenarios, the US Homeland Security Council (HSC) (21) uses an ICU proportion of 15% for the overall pandemic and 25.7% for the peak week. For seasonal influenza, CDC’s FluSurge 2.0 (22,23) assumes that a baseline of 15% of admitted influenza patients require ICU care; HHS makes similar assumptions (2) (online Technical Appendix Table 3). On the basis of these data and reports, we assumed peak-week ICU proportions of 20% during a mild pandemic and 25% during moderate and severe pandemics.

### Proportion of ICU Patients Requiring Ventilation

FluSurge 2.0 assumes 50% of patients with seasonal influenza admitted to the ICU require ventilation (22). HSC assumes 50% throughout a pandemic (21), and HHS uses 50.4% for a moderate scenario and 50% for a severe scenario (2) (online Technical Appendix Table 3). We assumed

that 50% of patients in the ICU who have pandemic influenza require ventilation across all scenarios.

### Proportion of Ventilated Patients Requiring 2 Weeks of Ventilation

FluSurge 2.0 (22) assumes that ventilatory support of ILI patients lasts 10 days. We have weekly time resolution and assumed 60% of patients receiving ventilatory support require only 1 week, and the remaining 40% require a second week.


### Simulating Pandemic Scenarios

We generated a mild scenario by fitting our forecasting model to hospital discharge data for the 2009 pandemic. Because comparable data are not available from 1957 and 1968 (moderate) and 1918 (severe), we scaled the 2009 estimates to model these scenarios. HHS (2) and HSC (21) use similar pandemic scaling factors, except HSC rates for hospitalization, ICU care, and mechanical ventilation are  $\approx 17\%$  and  $14\%$  lower than HHS rates for moderate and severe scenarios, respectively. (See [24] for scaling methods for an emerging pandemic.) CDC’s median estimate of hospitalizations for influenza A(H1N1)pdm09 (April 2009–April 2010) is 275,000. Combining this with the HHS scenario (online Technical Appendix Table 3), we scaled our mild pandemic hospitalization estimates by  $865,000/275,000 = 3.14$  and  $9,900,000/275,000 = 36$  to model moderate and severe scenarios, while preserving the variability, spatial correlation, and temporal correlation estimated for 2009.

### Results

Under the mild pandemic scenario, recommended stockpiles ranged from 200 to 400 ventilators (Figure 2, panel A). For example, if we specify the risk tolerance to be an EUD of at most 5 patients, then the recommended stockpile is 272 ventilators, including a central stockpile of 12. The PUD for this scenario, which is computed post hoc, is 30% (Figure 2, panel B). Thus, if the public health department builds the recommended central and local stockpiles, it can expect that no more than 5 patients statewide will go without ventilation, and a 70% chance exists that no demand anywhere will be unmet. As the risk tolerance decreases from an EUD of 5, the recommended stockpile grows sharply; as the EUD increases, the stockpile decreases nearly linearly (Figure 2, panel A). Ventilators are allocated primarily to local sites rather than to the central stockpile (Figure 2, panel C).

The optimal stockpile allocations under moderate and severe pandemic scenarios are qualitatively, but not quantitatively, similar (Figure 3). With an EUD tolerance of 5 patients, the recommended stockpiles increase to 1,172 and 15,697 ventilators for moderate and severe scenarios,



**Figure 2.** Optimal ventilator stockpiles for a mild pandemic scenario, Texas, USA. The total size of the optimal stockpile, summed across the central and 8 HSR stockpiles, decreases as risk tolerance increases. Risk for unmet demand for ventilators is quantified as the expected number of hospitalized influenza patients statewide not receiving necessary ventilation (EUD) (A) and the probability of at least 1 hospitalized patient in Texas not receiving necessary ventilation (PUD) (B). We optimized directly for EUD and calculated PUD post hoc. Red circles indicate EUD/PUD of 5 patients. C) Optimal allocation among central and regional sites when EUD is set to 5 patients, equivalent to a stockpile of 272 ventilators. EUD, expected unmet demand; PUD, probability of unmet demand; HSR, health service region.

respectively. These stockpiles scale roughly according to our assumptions that moderate and severe pandemics have hospitalization rates of 3.14 and 36 times higher than the mild pandemic, respectively, and that the fraction of hospitalized patients requiring ICU admission increases from 20% in the mild scenario to 25% in the other scenarios. Specifically, peak ventilator demand increases by factors of  $(0.25/0.20) \times 3.14 = 3.93$  and  $(0.25/0.20) \times 36 = 45$  from the mild to moderate and severe scenarios, respectively. This scaling would exactly predict how stockpiles would grow if we increased the risk tolerance by factors of 3.93 and 45. However, we fixed the EUD limit to 5 patients, so stockpile growth exceeds these scaling factors.

### Sensitivity Analysis

We assessed the sensitivity of the recommended stockpiling strategies to several factors. For a fixed risk tolerance (EUD), increasing the proportions of hospitalized patients requiring ICU admission and ventilation results in comparable increases in the recommended stockpiles. However, increasing the proportion of patients requiring 2 weeks of ventilation (rather than just 1) produces a slightly more complicated effect. Because the demand at peak week will depend on both established and newly admitted patients, increasing the 2-week proportion from 0 to 1 might not exactly double the demand. Based on 2009 pandemic hospitalization data, peak-week mean demand across Texas is expected to increase by a factor of 1.42 when the 2-week proportion increases from 0.4 to 1. The recommended stockpile grows accordingly. Under the mild pandemic scenario, the stockpile grows by a factor of 1.38 for an EUD near 0 ventilators and 1.42 for an EUD close to 5 ventilators.

We also varied the wastage rate for centrally held ventilators and the region-to-region correlation in peak demand. The baseline wastage of 0.2 means that 1 in 5

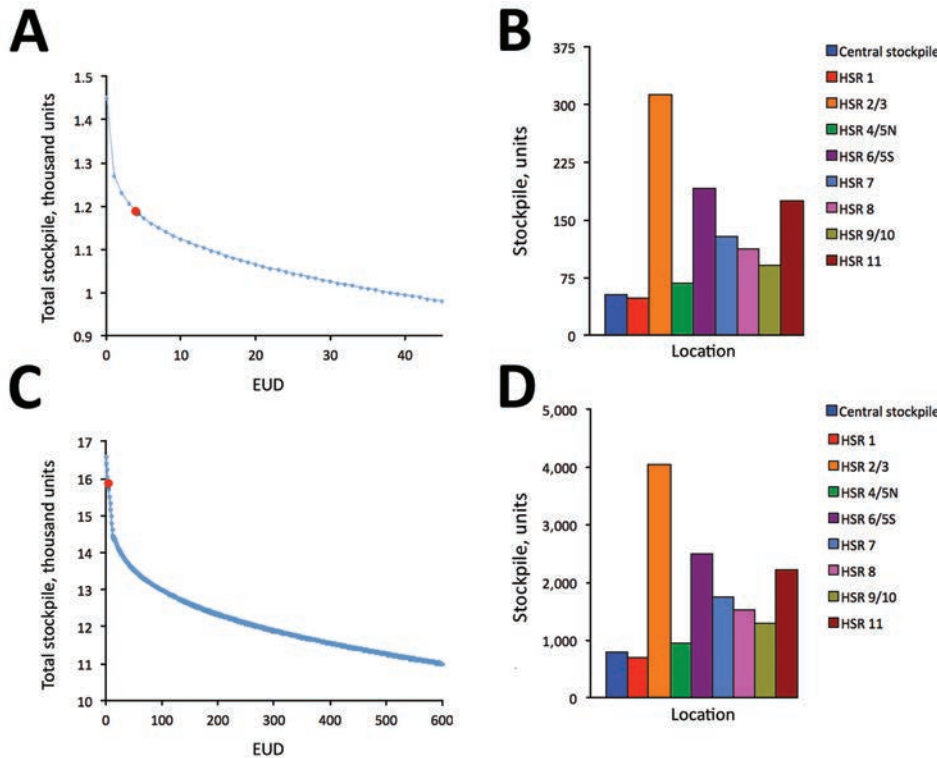
ventilators distributed from the central stockpile is not used effectively. This wastage contributes to relatively small recommended central stockpiles (e.g., just 4.4% of the total stockpile under the mild scenario with an EUD of 5 ventilators). As the wastage rate decreases, the central allocation slowly increases (Table 2; online Technical Appendix Figure 2). The benefit of risk pooling through a central stockpile also grows as the region-to-region correlation in peak demand shrinks (Table 2; online Technical Appendix Figure 3).

### Retrospective Analysis of 2009 Pandemic

During the 2009 pandemic, hospitals across Texas held an estimated 3,730 ventilators. When aggregated by region, the 8 HSRs had stockpiles ranging from 151 to 1,233 ventilators (online Technical Appendix Table 1). Under mild and moderate pandemic scenarios, we projected expected statewide demands for 230 and 903 ventilators, respectively, with each HSR holding a stockpile at least 6 SD above the forecasted mean demand. Given this ample regional surge capacity, there would have been no need for central stockpiling. Under the severe scenario, however, the projected statewide demand is 10,333 ventilators, far exceeding 2009 stockpiles.

### Discussion

Central stockpiles can save costs but are advisable only when spatial correlation in peak demand is sufficiently low and stockpile deployment is sufficiently reliable. Data from Texas suggest that influenza peaks strongly correlate across regions. Such synchrony undercuts the risk-pooling benefits of central stockpiles. Furthermore, successful deployment requires not only central maintenance and physical transportation of ventilators to patients in need, but also healthcare facilities and clinicians trained to administer and troubleshoot available ventilator models, which



**Figure 3.** Optimal ventilator stockpiles for moderate and severe pandemic scenarios, Texas, USA. The total size of the required stockpile, summed across the central and 8 HSR stockpiles, decreases as risk tolerance (EUD) increases, for both moderate (A) and severe (C) pandemic scenarios. For an EUD of 5 patients (red circles), total stockpiles would be 1,172 (A) and 15,697 (C); optimal allocations to central and regional stockpiles are shown for moderate (B) and severe (D) scenarios. EUD, expected unmet demand; PUD, probability of unmet demand; HSR, health service region.

might differ from those held locally. Pandemic-related staff absenteeism might exacerbate this challenge. Our model incorporates this limitation by assuming that fraction of stockpiled ventilators are wasted. When we considered a plausible wastage parameter of 20% (based on discussions with Texas DSHS about likely impediments to successful deployment), the model recommended that <10% of ventilators be held centrally.

The recommended allocations among central and local stockpiles hinge critically on the relative efficiencies of a local versus central stockpile, which are largely unknown and perhaps changing to favor central stockpiles as delivery technology continues to improve. We made the simplifying assumption that locally held ventilators are perfectly matched to patients, and we considered a range of potential wastage rates for centrally held ventilators. In general, the

more reliable central stockpile deployment, the more advisable a central stockpile. For example, assuming only 0.1% wastage, we found that that all ventilators should be held centrally, regardless of spatiotemporal correlations in peak demand (Table 2). Thus, as deployment and local capacities continue to improve, distance will become less of an issue, and the advantages of central stockpiles might outweigh their shortcomings.

Our surprisingly small central allocation stems from 2 additional factors. First, the uncertainty in our estimates of peak hospitalizations, based on 2009 pandemic data, is relatively low. Across Texas' 8 HSRs, the coefficient of variation (measuring the level of uncertainty) in peak demand for ventilators ranged from 0.17 to 0.36 and averaged 0.24 (online Technical Appendix Table 4). When we increase these coefficients governing uncertainty 3-fold, the

**Table 2.** Central stockpile size, as a percentage of total stockpile, as a function of wastage and region-to-region correlation in peak ventilator demand during an influenza pandemic, Texas, USA\*

Wastage, %	Central stockpile, %		
	Region-to-region correlation = 0.55	Region-to-region correlation = 0.70	Region-to-region correlation = 0.85
40	2.8	1.5	0.2
30	4.7	2.9	0.9
20	7.0	<b>4.4</b>	2.0
10	9.8	6.7	3.5
1	18.0	13.5	10.1
0.5	25.4	20.6	17.0
0.3	32.8	27.4	22.7
0.2	48.2	43.9	39.3
0.1	100	100	100

\*The results in this table are based on a mild pandemic scenario and a limit on expected unmet demand of 5 patients. Bold indicates the baseline value.

recommended central stockpile increases only from 4.4% to 10% of the total, assuming a mild pandemic and a risk tolerance (EUD) of 5 untreated patients. Second, the small central allocation depends on the risk tolerance. As the risk tolerance shrinks from an EUD of 5 patients, both the number of ventilators in the total stockpile and the percentage held centrally grow (online Technical Appendix Figure 3). Still, even at tighter risk tolerances and a smaller region-to-region correlation in peak demand of 0.55, the central stockpile is <10%.

Our retrospective analysis of the 2009 influenza A(H1N1) pandemic in Texas suggests that hospitals had enough ventilators on hand to treat all patients requiring mechanical ventilation throughout the pandemic. Although these quantities are expected to suffice for a moderate (1957- and 1968-like) pandemic, in which hospitalization rates roughly triple, they would fall far short in a severe (1918-like) pandemic. If we optimistically assume perfect deployment, that is, 0 wastage, by assuming timely delivery, adequately trained and available staff (respiratory therapists, nurses, and physicians), sufficient space to care for a potentially large number of patients, and requisite ancillary equipment and supplies, then even a central stockpile of 8,900 ventilators in Texas—the total number of SNS ventilators in 2010 (9)—would fall short, with an expected unmet demand of 576 patients.

This study was supported by the National Institutes of Health (Models of Infectious Disease Agent Study grant U01 GM087719-01) and CDC (Public Health Emergency Preparedness).

Dr. Huang is a senior data scientist at Precima, LoyaltyOne US, Inc., Chicago, Illinois, USA. He conducted this research while a PhD student in the Graduate Program in Operations Research and Industrial Engineering at the University of Texas at Austin. His research interests include optimization, public health, and supply chain management.

## References

1. US Department of Health and Human Services. Pandemic flu history. 2016 [cited 2016 Jun 16]. <http://www.flu.gov/pandemic/history/index.html>
2. US Department of Health and Human Services. HHS pandemic influenza plan. 2005 [cited 2016 Jun 16]. <http://www.flu.gov/planning-preparedness/federal/hhspandemicinfluenzaplan.pdf>
3. Sutton J, Tierney K. Disaster preparedness: concepts, guidance, and research. Fritz Institute Assessing Disaster Preparedness Conference; 2006 Nov 3–4; Sebastopol, CA, USA; 2006 [cited 2016 Jun 16]. <http://www.fritzinstitute.org/pdfs/whitepaper/disasterpreparedness-concepts.pdf>
4. Smetanin P, Stiff D, Kumar A, Kobak P, Zarychanski R, Simonsen N, et al. Potential intensive care unit ventilator demand/capacity mismatch due to novel swine-origin H1N1 in Canada. *Can J Infect Dis Med Microbiol*. 2009;20:e115–23. <http://dx.doi.org/10.1155/2009/808209>
5. Stiff D, Kumar A, Kisson N, Fowler R, Jovet P, Skippin P, et al. Potential pediatric intensive care unit demand/capacity mismatch due to novel pH1N1 in Canada. *Pediatr Crit Care Med*. 2011;12:e51–7. <http://dx.doi.org/10.1097/PCC.0b013e3181e2a4fe>
6. Ercole A, Taylor BL, Rhodes A, Menon DK. Modelling the impact of an influenza A/H1N1 pandemic on critical care demand from early pathogenicity data: the case for sentinel reporting. *Anaesthesia*. 2009;64:937–41. <http://dx.doi.org/10.1111/j.1365-2044.2009.06070.x>
7. Centers for Disease Control and Prevention. Strategic National Stockpile. 2015 [cited 2016 Jun 16]. <http://www.cdc.gov/phpr/stockpile/stockpile.htm>
8. Malatino EM. Strategic National Stockpile: overview and ventilator assets. *Respir Care*. 2008;53:91–5, discussion 95.
9. Jamieson DB, Biddison ELD. Disaster planning for the intensive care unit: a critical framework. In: Scales DC, Rubenfeld GD, editors. *The organization of critical care*. New York: Springer; 2014. p. 261–75.
10. American Association for Respiratory Care. Guidelines for acquisition of ventilators to meet demands for pandemic flu and mass casualty incidents. 2008 [cited 2016 Jun 16]. [https://c.aarc.org/resources/vent\\_guidelines\\_08.pdf](https://c.aarc.org/resources/vent_guidelines_08.pdf)
11. American Association for Respiratory Care. The Strategic National Stockpile ventilator training program. 2016 [cited 2016 Jun 16]. <https://www.aarc.org/resources/clinical-resources/strategic-national-stockpile-ventilator-training-program/>
12. Meltzer MI, Patel A, Ajao A, Nystrom SV, Koonin LM. Estimates of the demand for mechanical ventilation in the United States during an influenza pandemic. *Clin Infect Dis*. 2015;60(Suppl 1):S52–7. <http://dx.doi.org/10.1093/cid/civ089>
13. Ajao A, Nystrom SV, Koonin LM, Patel A, Howell DR, Baccam P, et al. Assessing the capacity of the US health care system to use additional mechanical ventilators during a large-scale public health emergency. *Disaster Med Public Health Prep*. 2015;9:634–41. <http://dx.doi.org/10.1017/dmp.2015.105>
14. Zaza S, Koonin LM, Ajao A, Nystrom SV, Branson R, Patel A, et al. A conceptual framework for allocation of federally stockpiled ventilators during large-scale public health emergencies. *Health Secur*. 2016;14:1–6. <http://dx.doi.org/10.1089/hs.2015.0043>
15. Timbie JW, Ringel JS, Fox DS, Pillemer F, Waxman DA, Moore M, et al. Systematic review of strategies to manage and allocate scarce resources during mass casualty events. *Ann Emerg Med*. 2013;61:677–689.e101. <http://dx.doi.org/10.1016/j.annemergmed.2013.02.005>
16. Wilgis J. Strategies for providing mechanical ventilation in a mass casualty incident: distribution versus stockpiling. *Respir Care*. 2008;53:96–100, discussion 100–3.
17. Texas Department of State Health Services and The University of Texas at Austin. Texas pandemic flu toolkit. 2013 [cited 2016 Jun 16]. <http://flu.tacc.utexas.edu>
18. Centers for Disease Control and Prevention. CDC estimates of 2009 H1N1 influenza cases, hospitalizations and deaths in the United States, April 2009 through January 16, 2010 by age group. 2010 [cited 2016 Jun 16]. [http://www.cdc.gov/h1n1flu/estimates/April\\_January\\_16.htm](http://www.cdc.gov/h1n1flu/estimates/April_January_16.htm)
19. Texas Department of State Health Services. Texas aggregate surveillance summary—novel influenza A H1N1, week ending 12/26/09. 2009 [cited 2016 Jun 16]. <http://www.dshs.state.tx.us/txflu/TX-cumulative-age-archive.shtm>
20. Petris G, Petrone S, Campagnoli P. *Dynamic linear models with R*. New York: Springer; 2009.
21. US Homeland Security Council. National planning scenarios version 21.3 final draft. 2006 [cited 2016 Jun 16]. <http://publicintelligence.net/national-planning-scenarios-version-21-3-2006-final-draft>
22. Zhang X, Meltzer MI, Wortley PM. *FluSurge 2.0: a manual to assist state and local public health officials and hospital*

administrators in estimating the impact of an influenza pandemic on hospital surge capacity (beta test version). 2005 [cited 2016 Jun 16]. [https://www.cdc.gov/flu/pandemic-resources/tools/downloads/flusurge2.0\\_manual\\_060705.pdf](https://www.cdc.gov/flu/pandemic-resources/tools/downloads/flusurge2.0_manual_060705.pdf)


23. Zhang X, Meltzer MI, Wortley PM. FluSurge—a tool to estimate demand for hospital services during the next pandemic influenza. *Med Decis Making*. 2006;26:617–23. <http://dx.doi.org/10.1177/0272989X06295359>

24. Reed C, Biggerstaff M, Finelli L, Koonin LM, Beauvais D, Uzicanin A, et al. Novel framework for assessing epidemiologic effects of influenza epidemics and pandemics. *Emerg Infect Dis*. 2013;19:85–91. <http://dx.doi.org/10.3201/eid1901.120124>

Address for correspondence: Hsin-Chan Huang, Operations Research and Industrial Engineering, The University of Texas at Austin, Austin, TX 78712, USA; email: [neo.huang@utexas.edu](mailto:neo.huang@utexas.edu)

## February 2015: Complicated Datasets

- Entry Screening for Infectious Diseases in Humans
- Timing of Influenza A(H5N1) in Poultry and Humans and Seasonal Influenza Activity Worldwide, 2004–2013
- Murine Typhus, Reunion, France, 2011–2013
- Evidence for *Elizabethkingia anophelis* Transmission from Mother to Infant, Hong Kong
- Microbiota that Affect Risk for Shigellosis in Children in Low-Income Countries
- pH Level as a Marker for Predicting Death among Patients with *Vibrio vulnificus* Infection, South Korea, 2000–2011
- Refining Historical Limits Method to Improve Disease Cluster Detection, New York City, New York, USA
- Naturally Acquired Antibodies against *Haemophilus influenzae* Type a in Aboriginal Adults, Canada
- Infectious Causes of Encephalitis and Meningoencephalitis in Thailand, 2003–2005
- Novel Reassortant Influenza A(H5N8) Viruses among Inoculated Domestic and Wild Ducks, South Korea, 2014
- Vesicular Stomatitis Virus–Based Vaccines against Lassa and Ebola Viruses
- Close Relationship between West Nile Virus from Turkey and Lineage 1 Strain from Central African Republic
- Ascariasis in Humans and Pigs on Small-Scale Farms, Maine, USA, 2010–2013
- Potentially Novel *Ehrlichia* Species in Horses, Nicaragua
- *Neisseria meningitidis* ST-11 Clonal Complex, Chile 2012
- Molecular Diagnosis of Cause of Anisakiasis in Humans, South Korea
- *Streptococcus suis* Infection in Hospitalized Patients, Nakhon Phanom Province, Thailand
- Exposure-Based Screening for Nipah Virus Encephalitis, Bangladesh
- Quantifying Reporting Timeliness to Improve Outbreak Control
- Tickborne Relapsing Fever, Bitterroot Valley, Montana, USA
- Simulation Study of the Effect of Influenza and Influenza Vaccination on Risk of Acquiring Guillain-Barré Syndrome
- *Lagenidium giganteum* Pathogenicity in Mammals
- Optimizing Distribution of Pandemic Influenza Antiviral Drugs
- Use of Insecticide-Treated House Screens to Reduce Infestations of Dengue Virus Vectors, Mexico
- Comparative Analysis of African Swine Fever Virus Genotypes and Serogroups
- Awareness and Support of Release of Genetically Modified “Sterile” Mosquitoes, Key West, Florida, USA
- Novel *Candidatus Rickettsia* Species Detected in Nostril Tick from Human, Gabon, 2014
- Outbreak of Henipavirus Infection, Philippines, 2014





# Stockpiling Ventilators for Influenza Pandemics

## Technical Appendix

### Materials and Methods 1: Forecasting Peak-Week Demand for Ventilators

We describe here the mathematical and technical details of the model used to forecast influenza-like-illness (ILI) hospitalizations. The purpose of the dynamic linear model (DLM) that we formulate is to generate accurate estimates of ILI hospitalizations as, and when, new information on the predictor variables becomes available. Formally embedding this “learning from experience” notion into the mathematical framework is one of the key merits of the Bayesian updating of the stochastic parameters in a DLM.

#### Forecasting of Hospitalizations

The predictors we use for forecasting include ILINet weekly reports for the state of Texas, and 4 time-indicator variables to account for the seasonality effect on ILI hospitalizations. We group months as September–October (S–O), November–December (N–D), January–February (J–F), and March–April (M–A). So, the corresponding indicator variable takes value 1 or all indicator variables are 0 for May–August. We considered models that also included predictors of school calendars, a humidity index, and Google Flu Trends, but for the significant look-ahead period we require for stockpiling ventilators these variables did not add significant predictive power to the model. Before proceeding with the details of the forecasting model for ILI hospitalizations, we specify notation.

#### Notation

$h_t$ : dependent variable of ILI hospitalizations at time  $t$  (weeks)

$z_t$ : independent variable of ILINet weekly reports at time  $t$

$\gamma_t^i$ : time indicators for season  $i$  at time  $t$ ,  $i \in I = \{1, 2, 3, 4\}$

$$\gamma_t^1 = \begin{cases} 1 & \text{if } t \text{ is in S-O} \\ 0 & \text{otherwise} \end{cases} \quad \gamma_t^2 = \begin{cases} 1 & \text{if } t \text{ is in N-D} \\ 0 & \text{otherwise} \end{cases}$$

$$\gamma_t^3 = \begin{cases} 1 & \text{if } t \text{ is in J-F} \\ 0 & \text{otherwise} \end{cases} \quad \gamma_t^4 = \begin{cases} 1 & \text{if } t \text{ is in M-A} \\ 0 & \text{otherwise} \end{cases}$$

We could formulate a static multiple linear regression model to study ILI hospitalizations in the following manner:

$$h = \beta^0 + \beta z + \sum_{i \in I} \alpha^i \gamma^i. \quad (1)$$

However, to incorporate the evolution of the predictors over time, which has significant importance in forecasting of ILI hospitalizations, we instead posit a dynamic linear regression model:

$$h_t = \beta_t^0 + \beta_t z_t + \sum_{i \in I} \alpha_t^i \gamma_t^i. \quad (2)$$

The critical difference between equations (1) and (2) is that the regression parameters are no longer static, evidenced by introducing the time subscript  $t$  in equation (2). The estimation of the random parameters in equation (2) can be performed recursively using the Kalman filter (1).

Let  $(H_t)_{t \geq 1}$  be the time series of ILI hospitalizations influenced by the nonrandom regression parameters corresponding to the independent variables; i.e., the regression coefficients for ILINet reports and the 4 time-indicators. The independent variables form the regression vector  $F_t$  at

time  $t$  while their coefficients are represented by the state vector  $\theta_t$ . The state matrix  $G_t$  is the evolution of the state vector through time. By introducing Gaussian measurement errors,  $v_t$ , and

Gaussian state evolution errors,  $w_t$ , the dynamic linear model is given by:

$$H_t = F_t^T \theta_t + v_t, \quad v_t \sim N(0, V_t)$$

$$\theta_t = G_t \theta_{t-1} + w_t, \quad w_t \sim N(0, W_t).$$

Here  $H_t$ ,  $v_t$ , and  $V_t$  are univariate, while  $F_t$ ,  $\theta_t$ , and  $w_t$  are  $p$ -dimensional vectors where, in our setting,  $p = 6$ , including the intercept term. The matrices  $G_t$  and  $W_t$  are  $p \times p$  in dimension.

## State Estimation and Observation Forecasting

The recursive procedure for updating the state vector  $\theta_t$  and forecasting the response variable  $H_t$  of the dynamic linear model now follows. At time  $t-1$ , for some mean  $m_{t-1}$  and covariance matrix  $C_{t-1}$  the information about the state  $\theta_{t-1}$  is presented with the posterior distribution:

$$(\theta_{t-1} | h_{1:t-1}) \sim N(m_{t-1}, C_{t-1}).$$

The recursive procedure starts at time 0 by choosing  $m_0$  and  $C_0$  to be the best guess regarding the mean and variance of the state vector. We use a subset of the data in the simple regression model (1) to construct the prior information about  $m_0$  and  $C_0$ . Through direct application of Bayes' theorem, we obtain that the prior distribution of  $\theta_t$  given  $h_{1:t-1}$  is Gaussian, i.e.,

$$(\theta_t | h_{1:t-1}) \sim N(a_t, R_t), \text{ with } a_t \text{ and } R_t \text{ being:}$$

$$a_t = E(\theta_t | h_{1:t-1}) = G_t m_{t-1}$$

$$R_t = \text{Var}(\theta_t | h_{1:t-1}) = G_t C_{t-1} G_t^T + W_t.$$

Next, the 1-step-ahead predictive distribution of  $H_t$  given  $h_{1:t-1}$  is also Gaussian, i.e.,

$$(H_t | h_{1:t-1}) \sim N(f_t, Q_t), \text{ with } f_t \text{ and } Q_t \text{ as follows:}$$

$$f_t = E(H_t | h_{1:t-1}) = F_t^T a_t$$

$$Q_t = \text{Var}(H_t | h_{1:t-1}) = F_t^T R_t F_t + V_t.$$

After obtaining the observation  $h_t$ , the filtering distribution of  $\theta_t$  is, again, Gaussian, i.e.,

$$(\theta_t | h_{1:t}) \sim N(m_t, C_t). \text{ The parameters } m_t \text{ and } C_t \text{ can be computed as follows:}$$

$$m_t = E(\theta_t | h_{1:t}) = a_t + R_t F_t e_t / Q_t$$

$$C_t = \text{Var}(\theta_t | h_{1:t}) = R_t - R_t F_t F_t^T R_t / Q_t,$$

where  $e_t = H_t - f_t$  is the forecast error. Our discussion here follows the book (1), which we refer to for further details.

### Multiple Steps Ahead Forecasting

For the purpose of producing demand scenarios for our optimization model, we must forecast hospitalizations and in turn, ventilator demand, many weeks into the future. Suppose we wish to forecast  $k$  weeks ahead. With the observed values of  $h_{1:t}$ , we can also forecast the future values of the state vector  $\theta_{t+k}$  and the observation  $H_{t+k}$ . Let  $a_t(0) = m_t$  and  $R_t(0) = C_t$ . Then, for  $k \geq 1$ , the distribution of  $\theta_{t+k}$  given  $h_{1:t}$  is Gaussian, i.e.,  $(\theta_{t+k} | h_{1:t}) \sim N(a_t(k), R_t(k))$ , with  $a_t(k)$  and  $R_t(k)$  being:

$$a_t(k) = E(\theta_{t+k} | h_{1:t}) = G_{t+k} a_t(k-1)$$

$$R_t(k) = \text{Var}(\theta_{t+k} | h_{1:t}) = G_{t+k} R_t(k-1) G_{t+k}^T + W_{t+k}.$$

The distribution of  $H_{t+k}$  given  $h_{1:t}$  is also Gaussian, i.e.,  $(H_{t+k} | h_{1:t}) \sim N(f_t(k), Q_t(k))$ , with  $f_t(k)$  and  $Q_t(k)$  as follows:

$$f_t(k) = E(H_{t+k} | h_{1:t}) = F_{t+k}^T a_t(k)$$

$$Q_t(k) = \text{Var}(H_{t+k} | h_{1:t}) = F_{t+k}^T R_t(k) F_{t+k} + V_{t+k}.$$

We use 1 year of historical seasonal influenza data to construct the prior for  $m_0$  and  $C_0$ , and we use 2009 pandemic data to fit the model and forecast  $k = 40$  weeks into the future.

### From Hospitalizations to Peak-Week Demand for Ventilators

We index the health service regions (HSRs) in Texas by  $r \in R$ . The DLM predicts hospitalizations on a weekly basis for each of the 8 HSRs in the form of a multivariate Gaussian distribution, providing the means ( $f_{r,t}$ ) and variances ( $Q_{r,t}$ ) for each region. We estimate the region-to-region correlations ( $\rho_{HSR}$ ) using historical data, and we assume this correlation to be identical for each pair of regions. To estimate the peak-week demand for ventilators from the forecasted hospitalizations, we employ 4 additional parameters: 1)  $p_i$ , the proportion of hospitalized ILI patients requiring ICU care; 2)  $p_v$ , the proportion of ICU patients requiring ventilation; 3)  $p_{tw}$ , the proportion of ventilated patients requiring 2 weeks of ventilation, under

the assumption that at most 2 weeks is needed; and, 4)  $\rho_{r,t}$ , 1-week lagged temporal correlation in ILI hospital admission in region  $r$  at time  $t$  generated by the DLM.

We calculate the mean weekly demand for ventilators in region  $r$  at time  $t$  as follows:

$$\mu_{r,t} = (p_{tw} f_{r,t-1} + f_{r,t}) p_i p_v. \quad (3)$$

We obtain the corresponding variance of weekly demand for ventilators, involving temporal correlation ( $\rho_{r,t}$ ), as follows:

$$\sigma_{r,t}^2 = Q_{r,t} p_i^2 p_v^2 + p_{tw}^2 Q_{r,t-1} p_i^2 p_v^2 + 2\rho_{r,t} p_{tw} \sqrt{Q_{r,t-1} p_i^2 p_v^2 Q_{r,t} p_i^2 p_v^2}.$$

We choose the peak-week demand in a region as the week with the largest mean according to equation (3). With the estimated region-to-region correlation ( $\rho_{HSR}$ ), we employ a standard Monte Carlo sampling algorithm (2) to generate independent and identically distributed (i.i.d.) samples of peak-week demand for ventilators as input to the optimization model, which we describe next.

## Materials and Methods 2: Optimization Model for Stockpiling

### Two-Stage Optimization Model

To optimize stockpiling decisions, we construct a 2-stage stochastic program. We index the regional sites by  $r \in R$ . The value of the central stockpile,  $x$ , and the value of the stockpiles at each site,  $s = (s_r)_{r \in R}$ , must be selected before observing the demand for ventilators

$d(\omega) = (d_r(\omega))_{r \in R}$ . The decision to ship ventilators from the central stockpile to site  $r$  is

captured by decision variable  $y_r(\omega)$ , and this decision is made after observing the demand

realization, indexed by  $\omega \in \Omega$ . In addition, if  $y_r(\omega)$  ventilators are shipped, then  $wy_r(\omega)$

represents the number of ventilators wasted so that only  $(1-w)y_r(\omega)$  ventilators can be used at

site  $r$ . Hence, the model seeks a balance between 1) the flexibility permitted by holding

ventilators centrally so that they can be distributed to where they are needed most, and 2) the fact

that locally held ventilators are more effective than those shipped from the central stockpile after

a pandemic begins. The optimization model for stockpiling is as follows:

$$\min_{x,s,y} x + \sum_{r \in R} s_r \quad (4a)$$

$$\text{s.t.} \quad \sum_{r \in R} y_r(\omega) \leq x, \forall \omega \in \Omega \quad (4b)$$

$$\mathbb{E}_\omega \left[ \sum_{r \in R} \left[ \left[ d_r(\omega) - s_r \right]^+ - (1-w)y_r(\omega) \right]^+ \right] \leq L \quad (4c)$$

$$x \geq 0, s_r \geq 0, y_r(\omega) \geq 0, \forall r \in R, \omega \in \Omega. \quad (4d)$$

The objective function we minimize in (4a) is the total stockpile of central and regional ventilators. Constraint (4b) says that the total number of ventilators distributed from the central stockpile to the sites cannot exceed the number of ventilators stockpiled centrally. We let

$\left[ d_r(\omega) - s_r \right]^+ = \max\{d_r(\omega) - s_r, 0\}$  represent the amount by which peak demand for ventilators exceeds the existing supply at site  $r$  under scenario  $\omega$ , and

$\sum_{r \in R} \left[ \left[ d_r(\omega) - s_r \right]^+ - (1-w)y_r(\omega) \right]^+$  represents the total shortfall of ventilators statewide after

distributing the central stockpile under scenario  $\omega$ . Thus constraint (4c) ensures that the expected shortfall of ventilators over all sites does not exceed the limit,  $L$ . Constraint (4d) enforces non-negativity for each decision variable. Note that  $d(\omega)$ ,  $w$ , and  $L$  are input data, and

$y(\omega) = (y_r(\omega))_{r \in R}$  are decision variables. By prespecifying the values of  $x$ ,  $s$ , or neither, 3 variations of the model can be formulated with respect to stockpiling decisions:

1. Given existing stockpiles at the regional sites, optimize the number of centrally held ventilators.
2. Given an existing central stockpile, optimize the number of ventilators at each site.
3. Jointly optimize the central and regional stockpiles, allowing us to assess the advantages of stockpiling ventilators centrally versus at the sites.

Model (4) is stated in the form of the third variation above, but the first 2 variations can also be handled by fixing decision variables  $s$  or  $x$ , respectively, to prespecified values.

We cannot solve model (4) directly for the following reasons. The summed shortfall of ventilators, i.e.,  $\sum_{r \in R} \left[ \left[ d_r(\omega) - s_r \right]^+ - (1-w)y_r(\omega) \right]^+$ , is a nonstandard random variable due to

the 2 positive-part operators within the summation, even though  $d(\omega)$  has the form of a multivariate normal distribution. More importantly, the decision variables,  $y(\omega)$ , representing shipments to the sites, adapt to the demand realization under scenario  $\omega$ , increasing the model's complexity. Hence, below we create a sampling-based variant of model (4), using a standard Monte Carlo sampling algorithm (2) to generate a set of i.i.d. samples of peak demands from the multivariate normal distribution we describe earlier.

### A Monte Carlo Approximation to the Optimization Model

Let  $i = 1, \dots, n$  index the sampled scenarios. Our sampling-based variant of model (4) is as follows:

$$\min_{x, s, y, u, v} x + \sum_{r \in R} s_r \quad (5a)$$

$$\text{s.t.} \quad \sum_{r \in R} y_r^i \leq x, \forall i = 1, 2, \dots, n \quad (5b)$$

$$u_r^i \geq d_r^i - s_r, \forall r \in R, i = 1, 2, \dots, n \quad (5c)$$

$$v_r^i \geq u_r^i - (1-w)y_r^i, \forall r \in R, i = 1, 2, \dots, n \quad (5d)$$

$$\frac{1}{n} \sum_{i=1}^n \sum_{r \in R} v_r^i \leq L \quad (5e)$$

$$x \geq 0, s_r \geq 0, y_r^i \geq 0, u_r^i \geq 0, v_r^i \geq 0, \forall r \in R, i = 1, 2, \dots, n. \quad (5f)$$

The objective function in (5a) is identical to that in (4a). Constraint (5b) is analogous to constraint (4b), where we add index  $i$  to variable  $y_r$  because shipments from the central stockpile to the sites occur after observing the demand realization. In constraint (5c),  $d^i = (d_r^i)_{r \in R}, i = 1, 2, \dots, n$ , are the samples of ventilator demands, and in constraint (5d),  $(1-w)$  is the proportion of centrally held ventilators dispatched to the site that can be used. These 2 constraints take care of the 2 positive-part operators in constraint (4c) by using 2 new decision variables,  $u_r^i$  and  $v_r^i$ . Given that these variables capture the positive parts, constraint (5e) is analogous to constraint (4c), and constraint (5f) again captures non-negativity of all decision variables. While we state models (4) and (5) for a fixed value of  $L$ , we view this as a bi-criteria model in which we can explore the tradeoff between the cost of the total stockpile (which we assume is proportional to the number of ventilators) and the limit on expected shortfall ( $L$ ).

## References

1. Petris G, Petrone S, Campagnoli P. Dynamic linear models with R. New York: Springer; 2009.
2. Devroye L. Non-uniform random variate generation. New York: Springer-Verlag; 1986.
3. US Department of Health and Human Services (HHS). HHS pandemic influenza plan. 2005 [cited 2016 Jun 16]. <http://www.flu.gov/planning-preparedness/federal/hhspandemicinfluenzaplan.pdf>
4. Texas Department of State Health Services. The health service regions. 2014 [cited 2016 Jun 16]. <http://www.dshs.state.tx.us/regions/state.shtm>

**Technical Appendix Table 1.** Existing regional stockpiles of ventilators in the state of Texas\*

Region	No. of existing ventilators
HSR 1	151
HSR 2/3	1,233
HSR 4/5N	247
HSR 6/5S	742
HSR 7	247
HSR 8	458
HSR 9/10	287
HSR 11	365

\*HSR, health service region.

**Technical Appendix Table 2.** Temporal correlation in the dynamic linear model between consecutive weeks, April–December 2009\*

Region	Minimum	Peak week	Median	Maximum
HSR 1	0.38	0.38	0.44	0.46
HSR 2/3	0.08	0.11	0.28	0.28
HSR 4/5N	0.19	0.19	0.23	0.24
HSR 6/5S	0.32	0.34	0.64	0.65
HSR 7	0.19	0.20	0.33	0.35
HSR 8	0.16	0.16	0.29	0.30
HSR 9/10	0.08	0.12	0.42	0.43
HSR 11	0.07	0.07	0.20	0.21

\*When the peak-week correlation is not the minimum correlation over the 9 months, the minimum instead occurs the week before the peak week. HSR, health service region.



**Technical Appendix Table 3.** Number of illnesses, healthcare utilization, and deaths associated with moderate and severe pandemic influenza scenarios\*

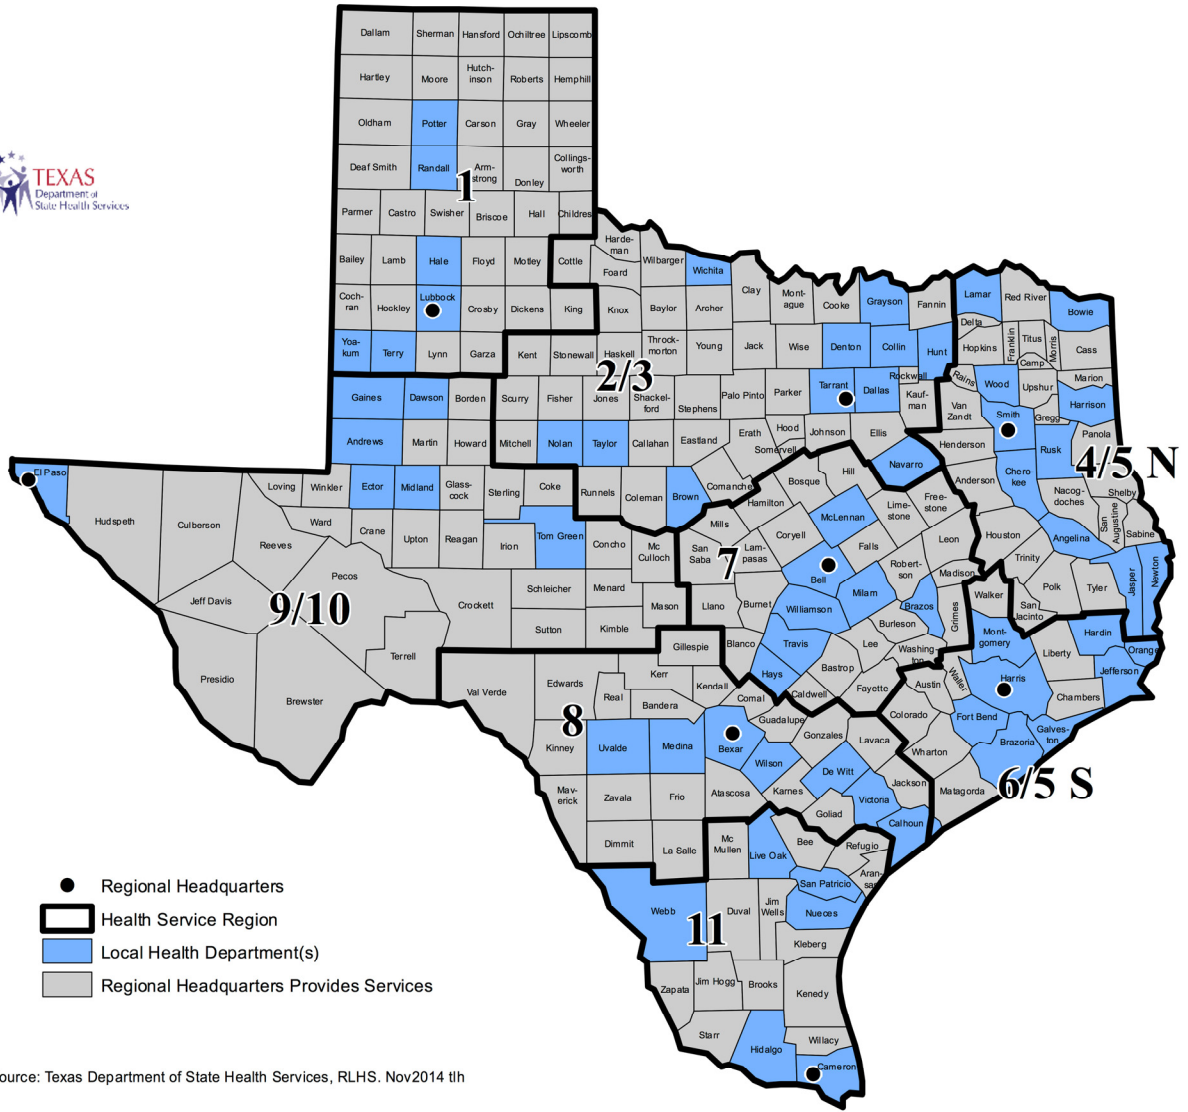
Characteristics	Moderate (1957- and 1968-like), no. (%)	Severe (1918-like), no. (%)
Illness	90 million (30)	90 million (30)
Outpatient medical care	45 million (50)	45 million (50)
Hospitalization	865,000	9,900,000
ICU care	128,750	1,485,000
Mechanical ventilation	64,875	742,500
Deaths	209,000	1,903,000

\*Source: (3). ICU, intensive care unit.

**Technical Appendix Table 4.** Estimated regional peak-week demand for ventilators in the mild scenario\*

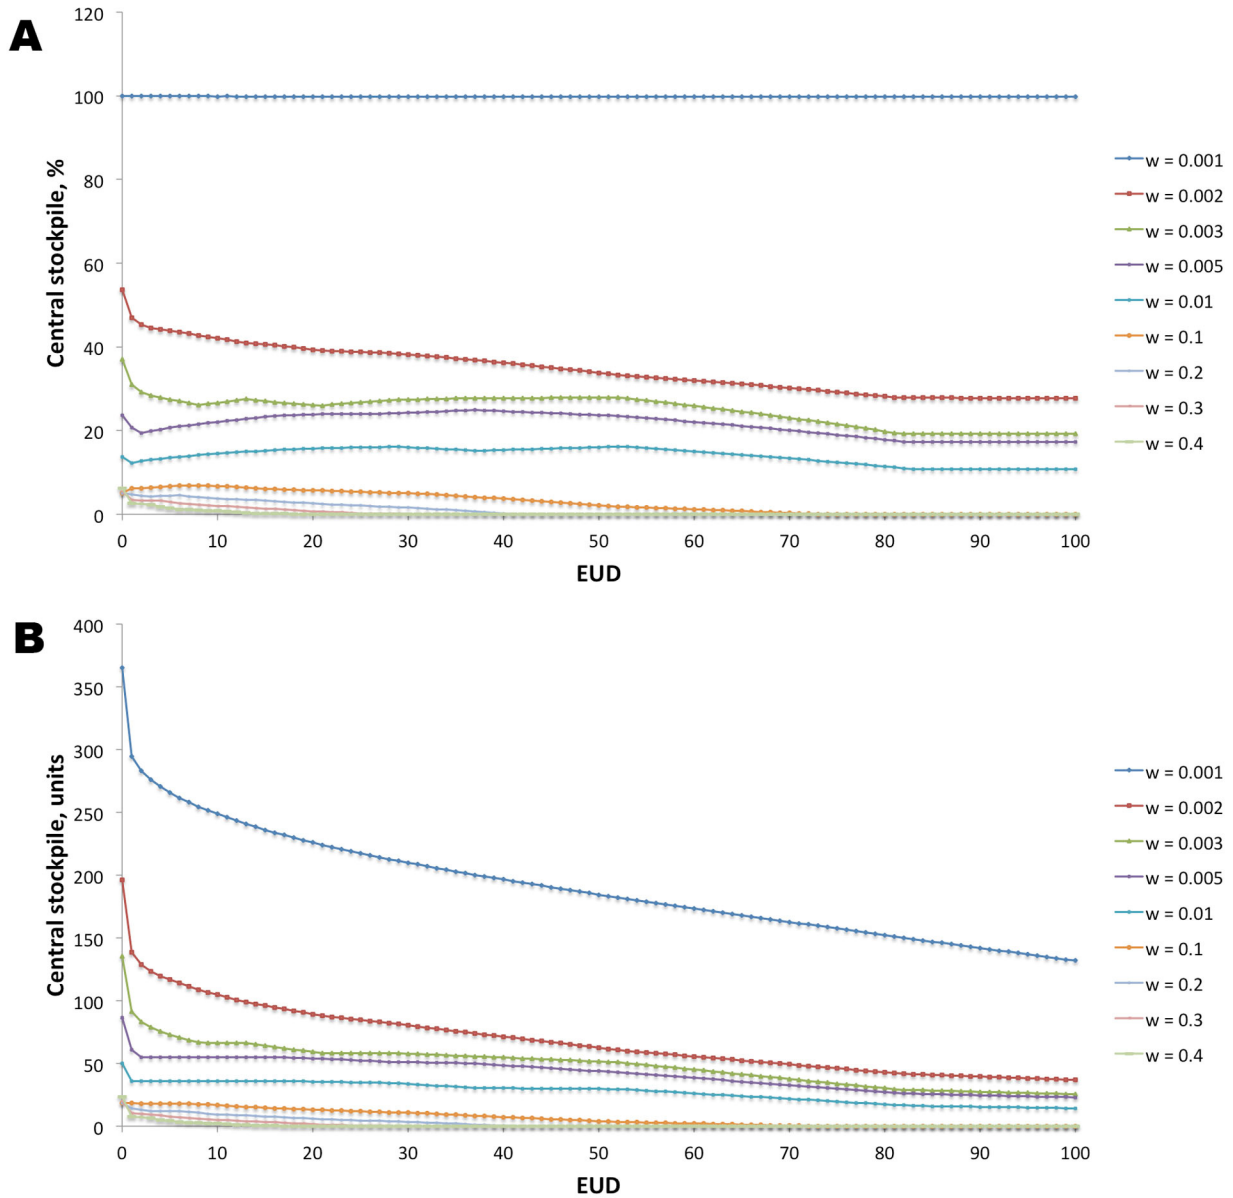
Region	Mean, units	± SD, units	Coefficient of variation
HSR 1	8.59	3.09	0.36
HSR 2/3	66.83	11.31	0.17
HSR 4/5N	12.93	3.48	0.27
HSR 6/5S	40.2	7.79	0.19
HSR 7	25.14	6.01	0.24
HSR 8	22.41	5.29	0.24
HSR 9/10	17.55	4.66	0.27
HSR 11	35.97	7.70	0.21

\*These estimates are based on April–December 2009 hospital discharge data in Texas. All the regional peak demands have a coefficient of variation <0.40, although the means range from 8.59 to 66.83. HSR, health service region.

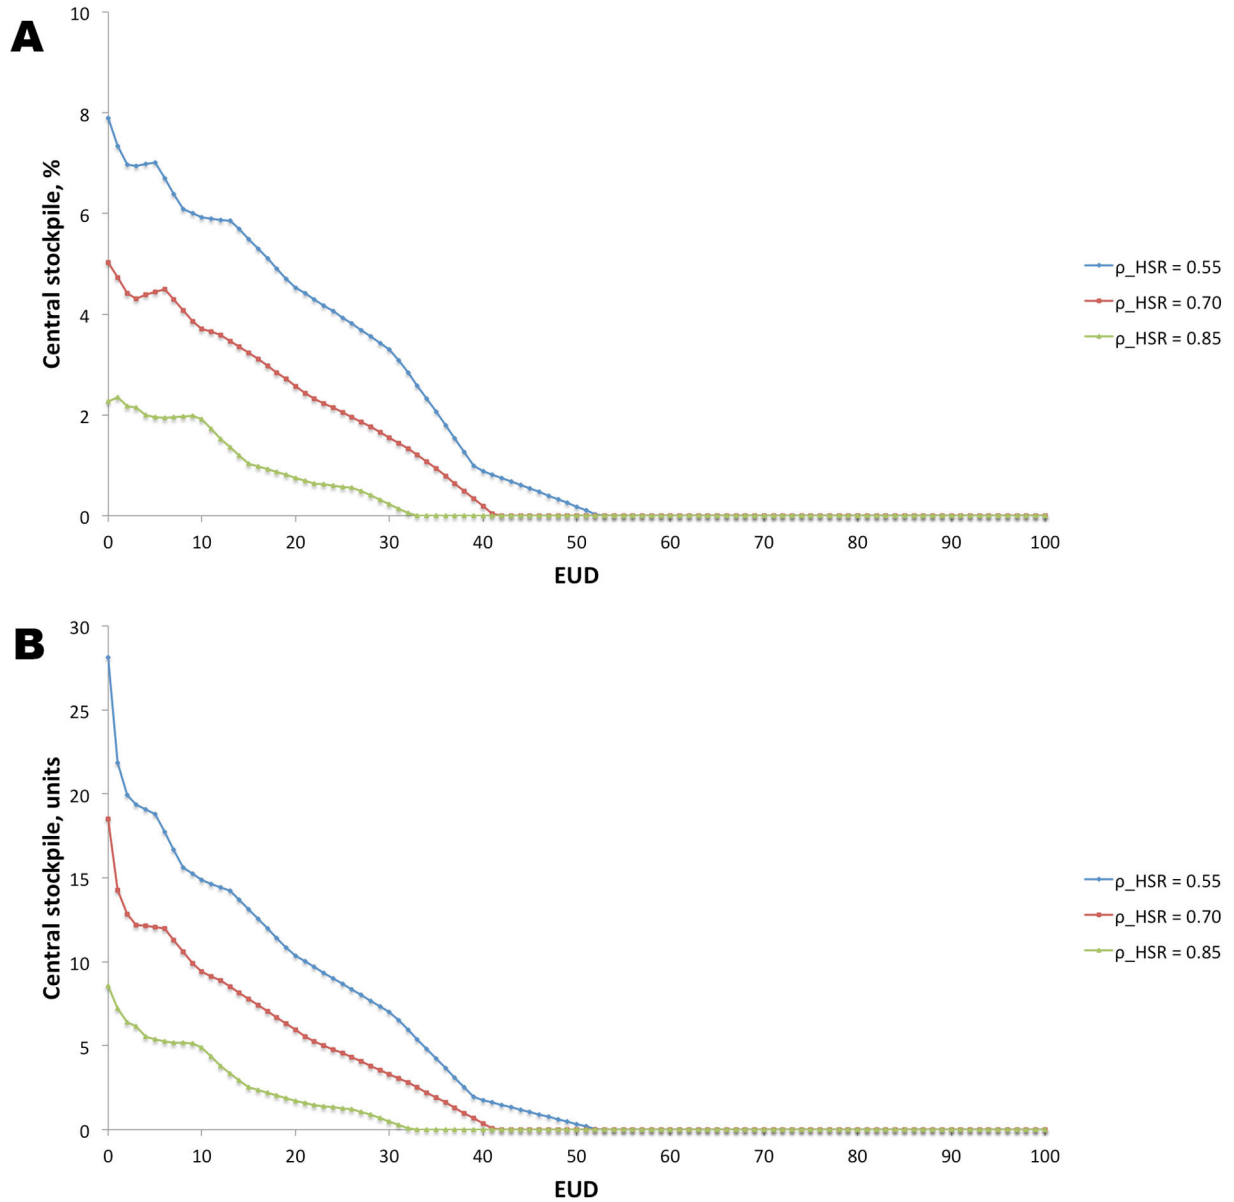


Source: Texas Department of State Health Services, RLHS. Nov2014 1th

Technical Appendix Figure 1. The 8 health service regions in Texas. Source: (4).



**Technical Appendix Figure 2.** The central stockpile versus EUD for various values of the wastage parameter ( $w$ ) for the mild influenza pandemic scenario, Texas, USA. The baseline result corresponds to  $w = 0.2$ , or 20%. A) Change in the percentage of the stockpile held centrally with the growth of EUD. B) Change in the number of ventilators held in the central stockpile. Mean peak-week demand, summed across all regions, is  $\approx 230$  ventilators in the mild scenario. EUD, expected unmet demand.



**Technical Appendix Figure 3.** The central stockpile versus EUD for various values of the region-to-region correlation coefficient ( $\rho_{HSR}$ ) under the mild influenza pandemic scenario, Texas, USA. The baseline result corresponds to  $\rho_{HSR} = 0.70$ . A) Change in the percentage of the stockpile held centrally with the growth of EUD. B) Change in the number of ventilators held in the central stockpile. EUD, expected unmet demand.