SANDIA REPORT
SAND2020-5277
Printed May 2020

Sandia National Laboratories

# Integrated Resource Supply-Demand-Routing Model for the COVID-19 Crisis

Christopher R. Frazier, Daniel J. Krofcheck, Jared L. Gearhart, Walter E. Beyeler

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico
87185 and Livermore,
California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.



Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods/

# ABSTRACT

As part of the Department of Energy response to the novel coronavirus disease (COVID-19) pandemic of 2020, a modeling effort was sponsored by the DOE Office of Science. Through this effort, an integrated planning framework was developed whose capabilities were demonstrated with the combination of a treatment resource demand model and an optimization model for routing supplies.

This report documents this framework and models, and an application involving ventilator demands and supplies in the continental United States. The goal of this application is to test the feasibility of implementing nationwide ventilator sharing in response to the COVID-19 crisis. Multiple scenarios were run using different combinations of forecasted and observed patient streams, and it is demonstrated that using a "worst-case" forecast for planning may be preferable to best mitigate supply-demand risks in an uncertain future. There is also a brief discussion of model uncertainty and its implications for the results.

## ACKNOWLEDGEMENTS

The authors acknowledge the support of the Office of Science for funding this work. We specifically acknowledge the following colleagues whose contributions assisted this work: Sean DeRosa, Patrick Finley, Katherine Klise, Laura Swiler, and Kate Tremba. We also acknowledge the generous support of our management, peers, and families as we took time from our regular duties and stay-at-home lives to perform this work under difficult and uncertain conditions during the COVID-19 Spring of 2020.

## CONTENTS



## LIST OF FIGURES

## LIST OF TABLES

This page left blank

## ACRONYMS AND DEFINITIONS

| Abbreviation | Definition |
|---|---|
| COVID-19 | Coronavirus Disease 2019 (formerly 2019 Novel Coronavirus (2019-nCoV)) |
| IHME | Institute for Health Metrics and Evaluation |
| ISDR | Integrated Supply-Demand-Routing Framework |
| JSON | JavaScript Object Notation |
| REST | Representational State Transfer |
| PPE | Personal Protection Equipment |

# 1. INTRODUCTION

As part of the Department of Energy response to the COVID-19 pandemic of 2020, a modeling effort was sponsored by the DOE Office of Science. Through this effort, multiple research and modeling efforts were launched, three of which are brought together for discussion and analysis in this report:

- A treatment demand model which provides predictions of medical resource needs for hospitals treating patients. The medical resources this model tracks include personnel (such as doctors and nurses), durable equipment (such as hospital beds and ventilators), and consumables (such as personal protection equipment (PPE) and medicine).
- An optimization model which determines routing paths for medical resources to match supply with demand. The model incorporates travel costs and seeks to minimize the number of regions with unmet demand.
- An integrated planning framework which has the capability of combining multiple supply, demand, and routing models into a single planning tool.

This report is focused on presenting the capabilities provided by combining these three together, focused specifically on an application looking at the feasibility of sharing ventilators across the U.S., and how best to plan for such a policy. It is important to emphasize that this capability is not intended to be used by policymakers to make actual resource routing decisions, but rather to provide insights into the patterns and scale of routing recommendations that arise under ensembles of uncertain future demand projections, to show the *feasibility* of specific routing strategies, and to understand the implications of making policy based on specific forecasts. Thus, the power of this work is not to provide specific directives for actions that need to be taken, but to provide strategic insight to assist policymakers in better appreciating the context and risks associated with a medical crisis.

Thus, this work should not be used to prescribe specific courses of action, but rather be used as a lens through which the uncertainty associated with resource demand consequences can be used to provide strategic logistic insight for policymakers in order to reduce time to critical decision making to mitigate life-threating resource shortfalls.

## 2. INTEGRATED SUPPLY-DEMAND-ROUTING FRAMEWORK

The integrated supply-demand-routing framework (ISDR) was developed to facilitate the combination of multiple healthcare resource models into a single framework. This capability provides a single interface for analysts and stakeholders to interact directly with real world data, providing valuable decision-making insights. Models are used to track resource supply and demand and routing models prescribe resource movements spatially and temporally to mitigate the occurrence of resource shortfalls. The healthcare resources tracked by this framework are plastic, and can be made to include human resources, (e.g., patients, healthcare workers), durable resources, (e.g., beds, ventilators), and consumables (e.g., PPE, medicines). The framework tracks resource supply and demand over time and consequently provides information on where and when there are expected to be resource surpluses and deficits, and how these are affected by resource movements.

Because ISDR is intended to be used as a tool for planning, internally it separates the concepts of forecasts and actualizations. Actualizations are the "reality" that the models experience - for example, the number of patients actually sent to hospitals. Forecasts are projections used internally by the models to make decisions that anticipate and are influenced by what is expected to happen. This allows planners to test scenarios where their forecasts are different from what actually happens. For example, a planner may make decisions based on forecasts of low patient counts, and if the actual patient counts are high, this framework would provide information on the consequences of such a mismatch. It is important to note that the distinction between actualizations and forecasts is internal to the framework, and that since typically ISDR is used to look into an unknown future, its internal actualizations will themselves be built from external forecasts; this subtlety is important for the subsequent discussion. Throughout this report, the terms actualization and observation will be used interchangeably, as will forecast and projection.

As ISDR steps through time, it communicates information between the various models and tracks their state. External resources are the primary means to drive ISDR through time. In its current use, these external resources are patients, but it is also possible that the resources could be additional healthcare supplies (such as hospital beds) being added to the system. Other elements that influence how ISDR steps through time are planned resource movements and scheduled status intervals. When ISDR steps forward to a new time period, all of the following steps will occur, in order, if they can:

- External resources (actualizations) are sent into ISDR and a router is used to determine where they are sent.
- Any planned resource movements are executed.
- The supply and demand models are updated based on the new and shifted resources, as well as the change in time.
- Forecasts are updated based on the observed external resources.
- The routing models use the updated forecast, supply, and demand information to plan immediate or future resource movements.

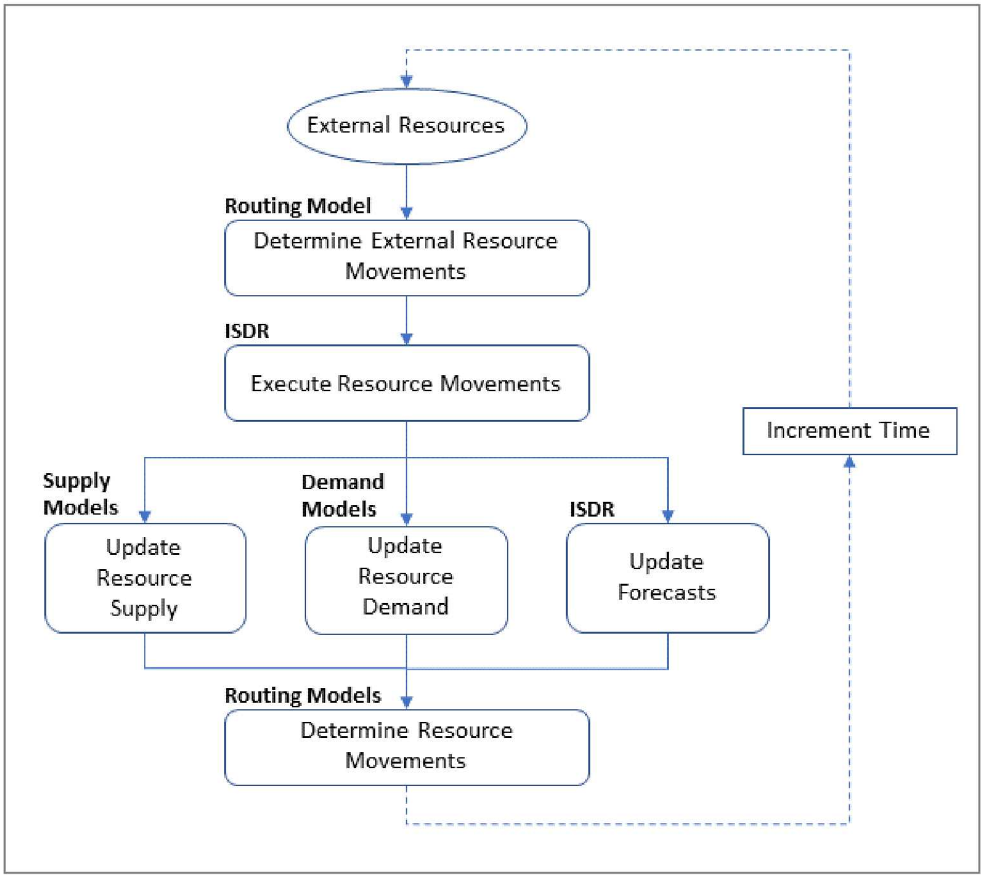


**Figure 2-1. Depiction of ISDR framework flow over time. At each step, multiple resources may be processed and tracked.**

An ISDR run completes when there are no more external resources to send into the system and all of the internal models have indicated that they have completed - currently this means that all patients have been treated (either they have died or been discharged).

The ISDR framework is implemented in the Java programming language. As part of the framework, generic components are available from which a specific ISDR model instance can be constructed. There is no requirement that any of the supply, demand, or routing models be implemented in Java, only that they be made accessible by the main Java framework.

# 3. TREATMENT DEMAND MODEL

The treatment demand model is used to determine the healthcare resources required to treat patients. It takes patients as inputs and, based on their characteristics, assigns them treatment paths which determine how long they will be treated and how many resources will be required. As time progresses, resources required for all of the patients flowing through the model are tracked, and at any given time the model can report the maximum amount of fixed or durable resources required, and the sum of the total consumables used.

Each patient has demographic attributes from which a set of possible treatment paths is determined along with probabilities, and from these a randomly selected treatment path is chosen. A specific treatment path has a series of stages, each of which requires a different number of resources. Additionally, there is uncertainty as to how long the treatment stage lasts, and whether the patient dies, is discharged, or continues on to the next treatment stage. Figure 3-1 depicts the flow of a single patient through the model. The model tracks multiple patients and resources simultaneously, continuing to run until all patients have completed their chosen treatment paths.

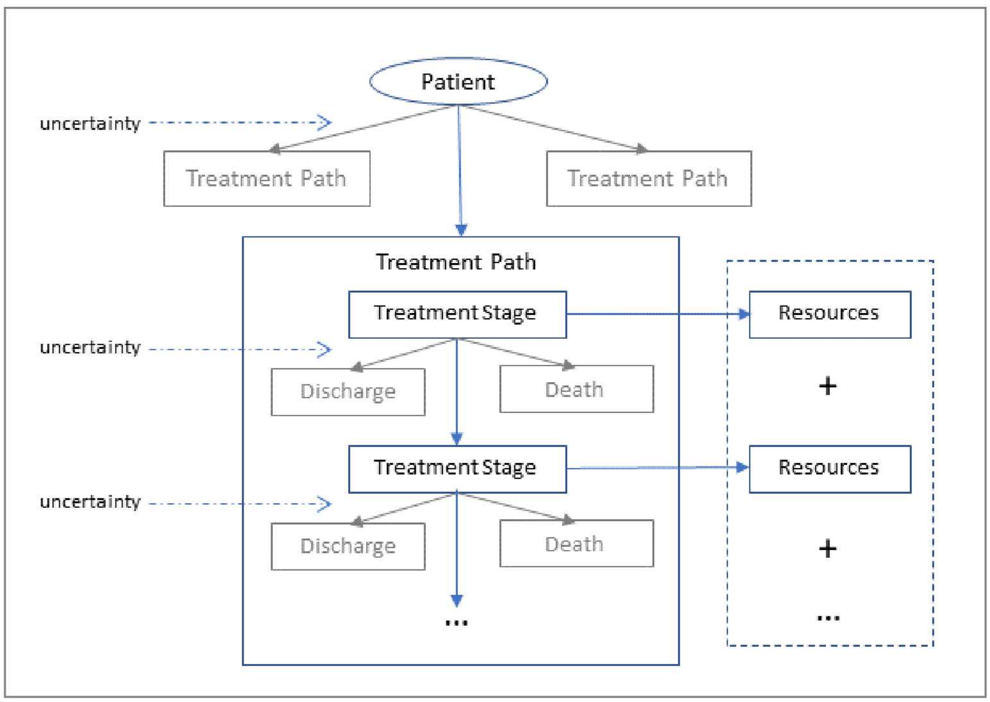


**Figure 3-1. A patient's flow through the treatment demand model, including selection of treatment path, progression through treatment stages, and tracking of resources.**

Example demographic characteristics for patients are age, sex, and whether they have preexisting health conditions such as hypertension or chronic obstructive pulmonary disease (COPD). Different treatment paths capture variations in disease symptoms, such as mild symptoms requiring short observation versus a severe case requiring ICU treatment and eventually a ventilator.

It is emphasized that there are several points of uncertainty as a patient moves through the model, but that the parameters themselves - the different probabilities, the resources used per treatment stage - are fixed. Because of this, the uncertainty captured in the model is focused on the variations in individual cases, but not on the model's structural uncertainties to its inputs. For an analysis focusing on this aspect of the model, refer to SAND2020-4900, *Uncertainty Analysis of Hospital Resource Demand Model for Covid-19*, which performs a thorough uncertainty quantification of the model, including its parameters [1].

The treatment demand model is implemented in the Java programming language. Typically, it is parameterized with input data saved in JavaScript Object Notation (JSON) format. It can be run from a representational state transfer (REST) interface, a command-line interface, and through direct access within a Java program.

# 4. ROUTING MODEL

Like the other components of the ISDR workflow, the routing model is a modular piece that can be modified, iterated on, or replaced entirely. In the current representation of the ISDR, we have implemented a single commodity optimization model, that aims to identify how a single resource should be moved between locations through time to minimize some resource shortfall. This formulation affords decision makers three key insights. First, because the routing model is responsible for intersecting demand estimates with resource inventory data, decision and policy makers can gain an understanding whether shortfalls may occur – and if so, this framework will provide descriptions of the timing, magnitude, and location of shortfalls to aid in logistic management efforts. Secondly, because the routing model makes recommendations for shuffling resources temporally and spatially to minimize shortfalls, existing collaborative relationships between groups (e.g., FEMA organizational units or medical consortiums) can be tested and described under a future of forecast uncertainty. Finally, this analysis framework affords stakeholders the understanding of the impact of forming new collaborative relationships between asset controllers that optimally mitigate resource shortfalls spatially and temporally.

This iteration of the routing model was intentionally designed to be lightweight, to increase the interpretability of the results and accelerate time to decision making in constrained planning horizons. The routing model takes as inputs a list of locations, a planning horizon, an initial inventory for the single commodity being analyzed at each location, a demand horizon for each location at each timestep in the planning horizon, and a pairwise cost associated with shipments for all locations being analyzed. The spatial and temporal resolution of the model are data driven, permitting investigations of shortfall mitigation strategies across states, nations, cities, hospitals, etc. This scale free nature of the routing model means the temporal and spatial resolution of the outputs are governed inherently by those of the inputs to the demand model.

The core model objective is to minimize the commodity shortfall at every time point and follows the formulation in eq1. The model achieves this optimization following several constraints and critical assumptions. Shortfalls between single commodity supply and demand are calculated at every time step – if supply exceeds demand, there is no shortfall (eq 2). Because the resource being tracked here is an essential component of life saving airway intervention in critical care environments, a shortfall is analogous to a patient mortality. As a result, shortfalls are not carried over from one timestep in the planning horizon to the next (eq. 3). The second constraint tracks inventory over time, and accounts for the specific availability of resources as a function of the forecast demand. Eq. 4 places a restriction on the values of the decision variables to be non-negative, and prevents locations from shipping resources to themselves.

$$\min \sum_{lt} x_{lt} \qquad \text{eq. 1}$$

$$\min \sum_{lt} x_{lt} \quad \text{eq. 1}$$

$$x_{\{lt\}} \geq d_{\{lt\}} - I_{\{lt\}} - \sum_{f} r_f x_{\{lt-f\}} \qquad \forall l, t \quad \text{eq. 2}$$

$$I_{\{lt\}} = I_{\{lt-1\}} + n_{\{lt\}} + \sum_{\{l't\}} s_{\{l'lt\}} - \sum_{\{l't\}} s_{\{ll't\}} \qquad \forall l, t \quad \text{eq. 3}$$

$$s_{ll't}, x_{lt}, I_{lt} \geq 0, \; s_{llt} = 0 \quad \text{eq. 4}$$

The above formulation represents the cost minimization of what is essentially 'excess demand' of the commodity, which is solely a function of our use case here but merits a brief discussion. As discussed above, this model was developed and applied in a context where the demand for some resource is being driven by the forecasted patient stream from an epidemiological model. The demand estimates here are the excess demand on the system, due to the patient stream being simulated. That is to say, the demand model will output the estimated *extra* demand on the system on top of background demand requirements. **Background demand estimates must be taken into account in order to not underestimate the potential commodity shortfall in critical planning contexts.** For example, in the application discussed in the following sections, we use pre-existing estimates of the percent utilization of ventilators on a state-by-state basis, and adjust the inventory data accordingly.

We further constrained the above model structure to better reflect real-world constraints on shipping and logistics management, aiming to consolidate redistribution shipments into as few bundles as possible. There is a fixed cost for opening a transportation link between pairwise combinations of locations, which minimizes coordinating pairs in the optimization solution. Eq. 5 includes the additional parameter $c^F_{ll'}$ fixed cost associated with making a shipment from location *l* to *l'* as an input, and a decision variable $f_{\{ll't\}}$ that indicates a shipment from location *l* to *l'* occurred at the start of time *t*.

$$\min \sum_{ll't,\, t \notin T} c^V_{ll'} s_{ll't} + \sum_{ll't,\, t \in T} C^{Disc} c^V_{ll'} s_{ll't} + \sum_{ll't,\, t \in T} c^F_{ll'} f_{ll't} \quad \text{eq. 5}$$

$$s_{ll't} \leq M f_{ll't} \qquad \forall l, l', t \quad \text{eq. 6}$$

$$f_{ll't} \in 0,1 \quad \text{eq. 7}$$

The outputs of the routing model include temporally explicit shortfalls, inventory data, and pairwise routing recommendations. These routing recommendations can then be construed as a graph stream of a temporal directed network, with the adjacency matric corresponding to establish shipment links, weighted by shipment magnitudes.

## 5. ISDR CASE STUDY EXAMPLE: VENTILATORS

As COVID-19 expanded to a pandemic, it became apparent that the most severely ill patients required ventilators to survive critical periods of the disease. In the United States, the initial extreme outbreaks of the disease were focused in specific regions which experienced shortages of many crucial healthcare resources given the rapidly increasing influx of patients; ventilators were among the most heavily scrutinized resources. Though there was uncertainty about the timing, magnitude, and geographic progression of the disease's spread, and therefore what the resource needs would be across the country, as the first flare-ups progressed it appeared that some regions would incur considerable shortages, whereas others would maintain surpluses of this critical resource.

As the crisis unfolded over March and April of 2020, one strategy that was raised by a number of people - including the governor of New York during his daily press appearances - was to encourage the collaboration of regions with surplus ventilators, facilitating resource reallocation to regions with current or expected shortfalls [2]. The proposed plan would require collaborators to track the progression of the disease and resource demands, dynamically adjusting inventories to mitigate shortfalls in other regions as new regions began to experience increases in COVID-19 cases. Putting aside concerns about politics, equity, and ventilator durability, it is worth investigating whether such a strategy is even feasible, and if it is, how to go about planning for it. We used this set of questions as the impetus to highlight the role that ISDR can play in communicating real-world solutions to logistics allocation optimization challenges in critical decision-making contexts. The following analysis uses ISDR with the treatment demand model and the routing optimization model to address these questions.

### 5.1. Ventilator Supply and Demand Models

Our analysis splits the U.S. into 49 regions: the 48 continental states plus Washington D.C. ISDR assigns each region its own treatment demand model, driven by patient flows to model ventilator demands. We parameterized the model using demographics based on the U.S. census and treatment paths based on published case studies, thus developing a set of parameters used for each of the treatment demand models. We chose to use forecasts from the University of Washington's Institute for Health Metrics and Evaluation (IHME) in this example. IHME develops state-level patient hospitalization forecasts at three levels: the lower uncertainty bound, the mean, and the upper uncertainty bound. The forecasts used in this analysis were published on April 7, 2020 [3]. The patient streams are at a daily resolution, and the framework runs its updates at the same frequency.

For the ventilator supplies, data has been aggregated from raw state to a US regional distribution. These provide initial ventilator counts for each region, and the analysis assumes that no new ventilators are introduced into the system over time.[1]

### 5.2. Patient Demand Forecasting

To forecast ventilator demands, each region has a separate treatment demand model driven by patient streams. The base patient streams used in these models are IHME forecasts, and these forecasts may be different from those used for the observed patient streams. For example, the observed patient stream for the ISDR model may be the IHME mean, but the ventilator forecasts

[1] While there was an effort to ramp up ventilator production - both new units and modifications of other devices - forecasts for these new ventilators were unavailable or vague. Also, not including these potential new supplies does not negate the general usefulness of this analysis as it stands. ISDR can readily accommodate production nodes to allow the creation of supplies, contingent on accurate input data.

could be driven instead by the IHME lower patient stream. Because these two may not be the same, at any point in time the observed ("real") patient streams may not equal the values used for the forecast. In a "real-world" situation, if the ground-truth did not match the projections, the projections would be adjusted to align with the ground-truth. To capture such behavior, the forecasted patient stream is scaled so that its value at the current model time matches that of the observed patient flow.

The scaling factor is calculated based on a comparison of the forecast and observed patient flows, and then every forecast point from the current model time is multiplied by this scaling factor and rounded to the nearest integer. The base scaling factor ($f_T$) is the ratio of the sum of each stream up to current model time (*T*):

$$f_T = \frac{\sum_{t=0}^{T} stream_t^{observed}}{\sum_{t=0}^{T} stream_t^{forecast}}$$

This scaling factor will run into problems if either the observed or forecasted patient stream sums are zero. If that is the case, different algorithms utilizing a broader time range must be used. If the observed patient stream sum is zero, then the ratio is estimated by supposing the forecast was zero up until the current model time. That is, the denominator is the sum of the forecast over its entire time range, and the numerator is the forecast sum minus the sum of the forecast to current model time:

$$f_T = \frac{\left(\sum_{t} stream_t^{forecast}\right) - \left(\sum_{t=0}^{T} stream_t^{forecast}\right)}{\sum_{t} stream_t^{forecast}}$$

If the forecast sum is zero, then the average forecast value per time over its entire time range ($\Delta t$) is calculated, and then a forecast denominator is estimated as this average multiplied by the total current model time that has elapsed:

$$f_T = \frac{\sum_{t=0}^{T} stream_t^{observed}}{\frac{T}{\Delta t}\left(\sum_{t} stream_t^{forecast}\right)}$$

If a forecast value does not exist at a specific model time, then an estimate (linear) is calculated from the nearest existent forecast points.

Note that the adjustment algorithm will not force the observed and forecasted patient streams to be equal at the current model time, and it is possible to have large discontinuities when the two diverge significantly from one another (Figure 5-1). This is reasonable as the forecast is intended to represent both the future progression of the outbreak (not its current state) and a general sentiment about it, and overfitting to the match the observed data can subvert this intent.[2]

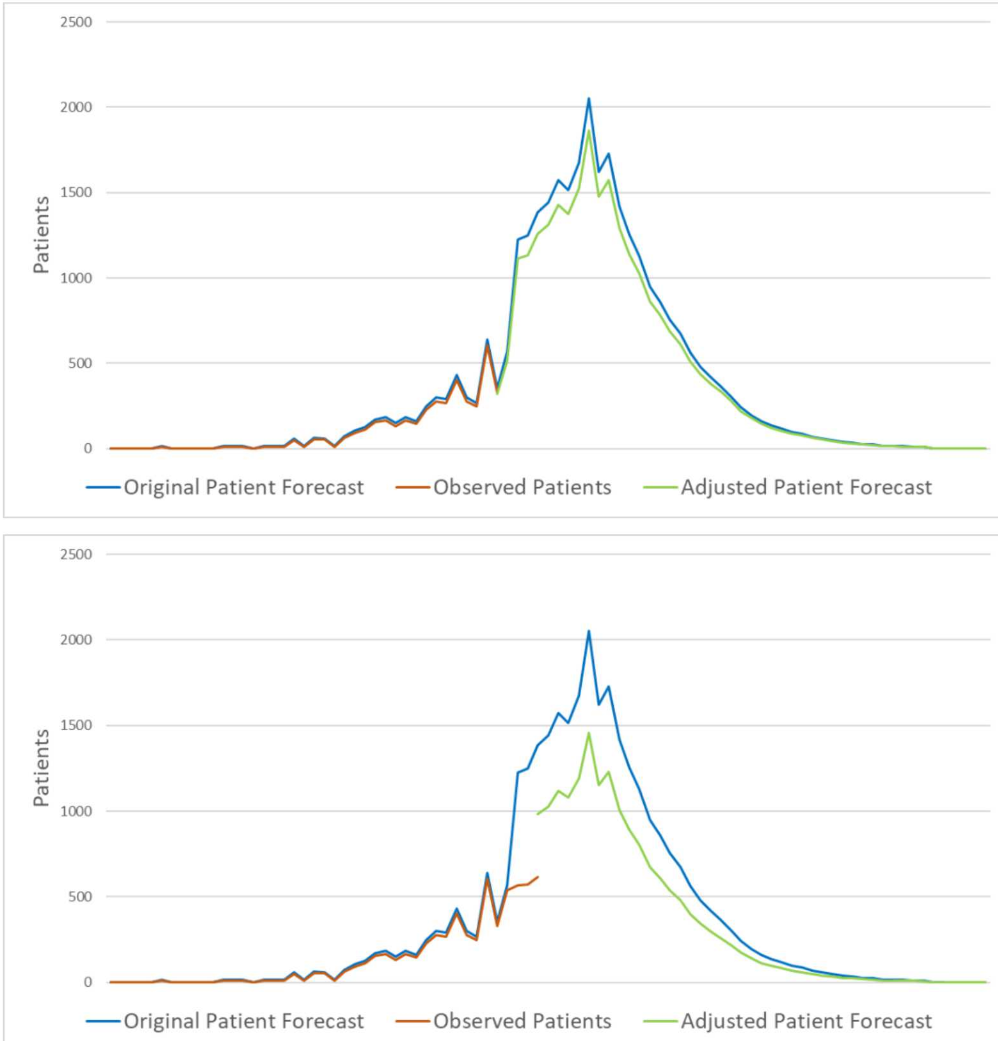


**Figure 5-1. Comparison of patient forecast adjustment when the observed stream is close to the projection (top) and when there is a significant separation between the two (bottom).**

## 5.3. Ventilator Routing Model

The routing model is set up to use the count of ventilators in each region at model time as the initial supply, and projected ventilator demands modeled from projected patient flows. Costs are calculated as the Euclidean distance between state centroids. The routing model is run every ten model days, uses forecasts for the next 30 days, and determines what resource movements to plan ten days in the future. By only looking one month into the future, the model is prevented from trying to

[2] It is possible to have a more sophisticated forecast adjustment algorithm which forces continuity at current model time but better preserves the forecast curve by incorporating the forecast's entire integral (similar to what is done if the observed sum is zero at model time). However, this increases the algorithmic complexity without improving the model performance, as the general shape of the forecasts is similar between the two methods; the point of continuity is not utilized and thus provides no real benefit.

incorporate information that may be too uncertain (especially given the forecast adjustment algorithm) and scheduling only ten days out allows it to be responsive to changes in observed conditions. These decisions were made to best emulate real world conditions, where epidemiological forecast uncertainty increases substantially with the forecast horizon.

# 6. ISDR CASE STUDY MODEL RESULTS

The ISDR ventilator model was run using different combinations of the IHME patient streams for forecasts and actualizations. The intent of this was to see what effect choosing different patient forecasts has on the feasibility of routing ventilator supplies to meet actual demands. As there are three IHME patient streams - low, mean, and upper - there are nine combinations of forecasts/actualizations, shown in Table 6-1 along with shorthand notations for them.

**Table 6-1. The nine combinations of observed and forecast patient streams using IMHE projections. The values in the center of the table will be used as shorthand references in the results discussion.**

| Observed Patient Streams | Projected Patient Streams | | |
|---|---|---|---|
| | **IMHE Lower** | **IMHE Mean** | **IMHE Upper** |
| **IMHE Lower** | lower/lower | lower/mean | lower/upper |
| **IMHE Mean** | mean/lower | mean/mean | mean/upper |
| **IMHE Upper** | upper/lower | upper/mean | upper/upper |

When performing the runs, the same random seed was used so that the probabilistic choices made in the treatment demand models were identical across all model runs. This allows comparisons between the runs to be due to the differences in the dynamic optimization solutions rather the uncertainty associated with stochastic choice. We further discuss the incorporation of model uncertainty below (Section 7), as well as the consequences of this single-seed simplification.

## 6.1. Aggregate Ventilator Metrics

The regular time step for each model run is 24 hours, and after every day is finished ISDR provides status information about ventilator supply and demand at that time. This affords the temporal description of a number of metrics:

- *Ventilator Supply* – This indicates where ventilators currently reside. Because the total supply is fixed (no new ventilators are introduced into the system during a model run), aggregating across the system results in a constant.
- *Ventilator Demand* – This indicates where ventilators are needed. Aggregating across the system gives the total number of ventilators needed for patients being treated at any given time. Because this is solely dependent on the treatment demand model results and because a fixed random seed is being used, the ventilator demand will be the same for scenarios using the same patient stream observations.
- *Ventilator Deficit (Shortfall)* – This is the number of ventilators needed that are not being supplied. Calculating by location and then aggregating across the system provides the total unmet demand for ventilators. In our routing model assumptions, this unmet deficit is not carried over from timepoint to timepoint, as ventilator demand is immediate and cannot be

delayed (the consequence of a ventilator shortfall is worse outcomes for the patients and increased chances of death).

- *Unused Ventilator Supply* – This is the number of ventilators that are not in use. Calculating by location and then aggregating across the system provides the total potential capacity for ventilator demand.
- *Shipment Cost* – The model also tracks the total cost associated with shipments, allowing the temporal description of the cost of achieving some shortfall mitigation target. We exclude these results in this case study however for two reasons. First, in this example we never allowed cost to be a reason to *not* ship a ventilator if the supply existed – doing so given our constraints would effectively place a cost on human life. Second, the use of simple Euclidian distance for the shipment cost is used to drive the model to prefer shorter movements, not to make economic tradeoffs.

Note that because ventilator demand may only be satisfied if the supply is in the same location, it is possible to have unused ventilator supply and still have a non-zero deficit. Such a situation indicates that the routing model was not able to adequately move supplies to meet demands, which could happen for two reasons: insufficient responsiveness due to the 10-day gap between ventilator movements, or a difference between the expectations provided by forecasts and the actualizations actually observed.

Table 6-2 and Table 6-3 presents the results of model runs, the former for results constant across runs using the same observed patient streams, the latter for the results which differ across the nine observation/forecast patient stream combinations. In Table 6-2, The "Minimum Possible Maximum Deficit" column is the smallest maximum deficit possible if the ventilator supply could just be magically moved around the country to satisfy demand. The "Maximum Possible Minimum Unused Supply" is similar, only here presenting the largest minimum unused supply. Both of these columns are calculated directly from the ventilator supply and maximum demand columns. In Table 6-3, the "Maximum Deficit" and "Minimum Unused Supply" columns provide the most extreme values of these metrics over the course of the model runs.

When the forecasts and actualizations align (the diagonals of Table 6-3), the routing model does a good job matching ventilator supplies to demands. For lower/lower and mean/mean, the maximum deficit is close to the theoretical best case (minimum possible maximum deficit) and is small compared with the maximum demand – 0.66% and 0.35%, respectively. For upper/upper, there is an insurmountable gap between with maximum demand and the fixed ventilator supply, however, the model run is able to get very close to the theoretical best case, deviating by only 0.2% from it.

Likewise, for all three cases, the difference between the minimum unused supply and the theoretical best case (maximum possible minimum unused supply) is far less than 1% of the total ventilator supply. This indicates that if the forecasted patient streams align with what is actually observed, then it is feasible to perform routing at 10-day intervals to meet demand, and that deviance from the optimal case is very small.

If the observed patient streams are lower than those projected – the lower/mean, lower/upper, and mean/upper scenarios – then the models still perform fairly well. In fact, in all three cases, the theoretical best possible unused supply is achieved, which outperforms even the models where the forecasts and observations matched. For the ventilator deficits, the lower/mean and lower/upper runs also were very close to or equaled the theoretical best case. However, for the mean/upper scenario, the maximum deficit of -1,136 is well off the theoretical best case of zero, being 7.1% of the maximum demand for this scenario. This indicates that while over-predicting patient streams generally performs well, it is not necessarily an ideal strategy, as it is shows that it may not be able to meet the necessary demand, even if there are enough ventilators to do so. It is important to note that in the mean/upper case, there are actually at least 12,497 *excess* ventilators, so the interaction of the observed/forecast patient stream difference with the routing model is causing clear inefficiencies in matching supply to demand.

**Table 6-2. Model run results which are constant across all runs for a given observed patient stream.**

| Observed Patient Stream | Ventilator Supply | Maximum Demand | Minimum Possible Maximum Deficit | Maximum Possible Minimum Unused Supply |
|---|---|---|---|---|
| IMHE Lower | 28472 | 7858 | 0 | 20614 |
| IMHE Mean | 28472 | 15975 | 0 | 12497 |
| IMHE Upper | 28472 | 45168 | -16696 | 0 |

**Table 6-3. Model run results aggregated across the system for the nine different observation/forecast patient stream combinations.**

| Observed Patient Stream | Maximum Deficit (by Forecast Patient Stream) | | | Minimum Unused Supply (by Forecast Patient Stream) | | |
|---|---|---|---|---|---|---|
| | IMHE Lower | IMHE Mean | IMHE Upper | IMHE Lower | IMHE Mean | IMHE Upper |
| IMHE Lower | -52 | 0 | -9 | 20643 | 20614 | 20614 |
| IMHE Mean | -3707 | -57 | -1136 | 16045 | 12541 | 12497 |
| IMHE Upper | -24721 | -22796 | -16723 | 7949 | 6025 | 9 |

When the forecasted streams under-predict the observed patient streams, the model runs perform significantly worse. The maximum deficit for the mean/lower run is over three times that of the mean/upper run just discussed, indicating a significant deviation from what is possible given the ventilator supply and maximum demand for the scenario. Likewise, the minimum unused supply is 28% higher than the theoretical best case, further indicating the problem of the models being unable to match supply to demand.

The upper/lower and upper/mean scenarios deviate even further from the theoretical ideals. While this may be partly due to the much higher demand generated by the scenario's observed patient stream, it appears that generally speaking, under-forecasting the observed patient streams leads to significantly worse outcomes. Also, there appears to be a trend within the data: the further the forecast is from the observed result (that is: upper/lower is further away than upper/mean), the metrics show poorer model performance.

The issues that appear to arise for using forecasts that are lower than the actualizations are shown graphically in Figure 6-1 and Figure 6-2, which show metric time-series charts for the runs realizing observations from the IMHE Mean and IMHE Upper patient streams, respectively. In both charts it is seen that the scenarios where a forecast less than the observed is used (mean/lower, upper/lower, and upper/mean), both the ventilator deficits and the unused ventilator supply are generally worse than the scenarios where the forecasts equal or exceed the observed patient streams. These charts show that this conclusion is not due to outlier points in the results, but rather general trends that are happening consistently over the entire time of the model run.

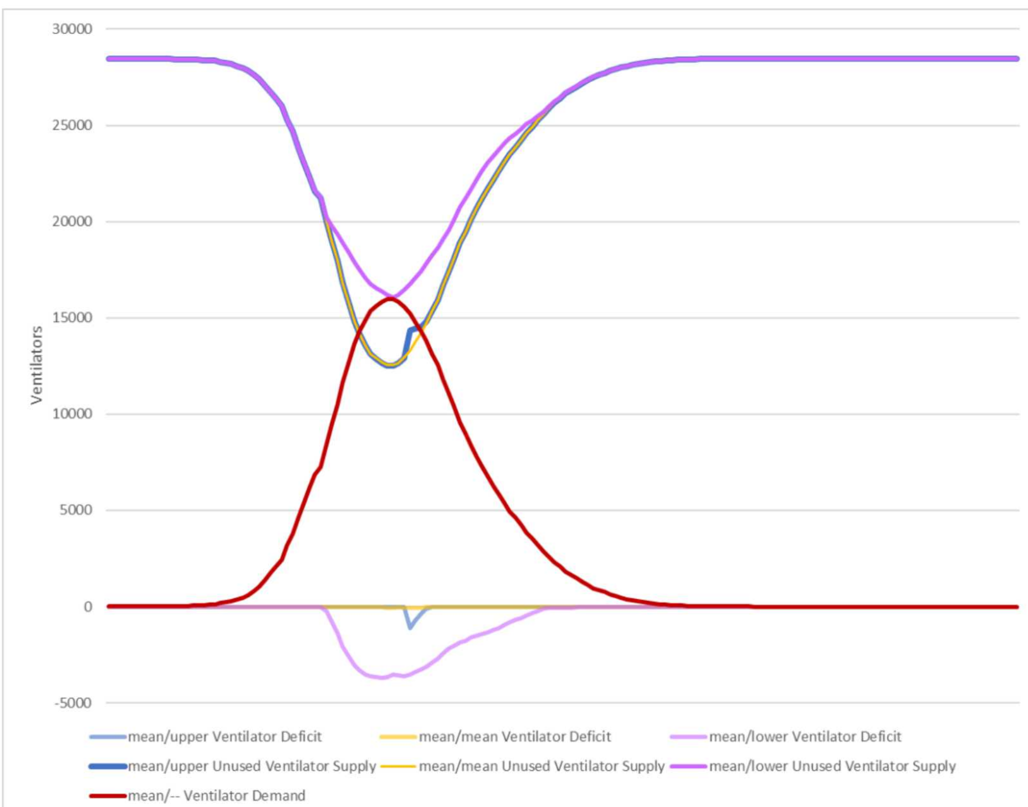


**Figure 6-1. Ventilator demand, deficit, and unused supply time-series for the scenarios using IHME Mean patient streams for observations.**

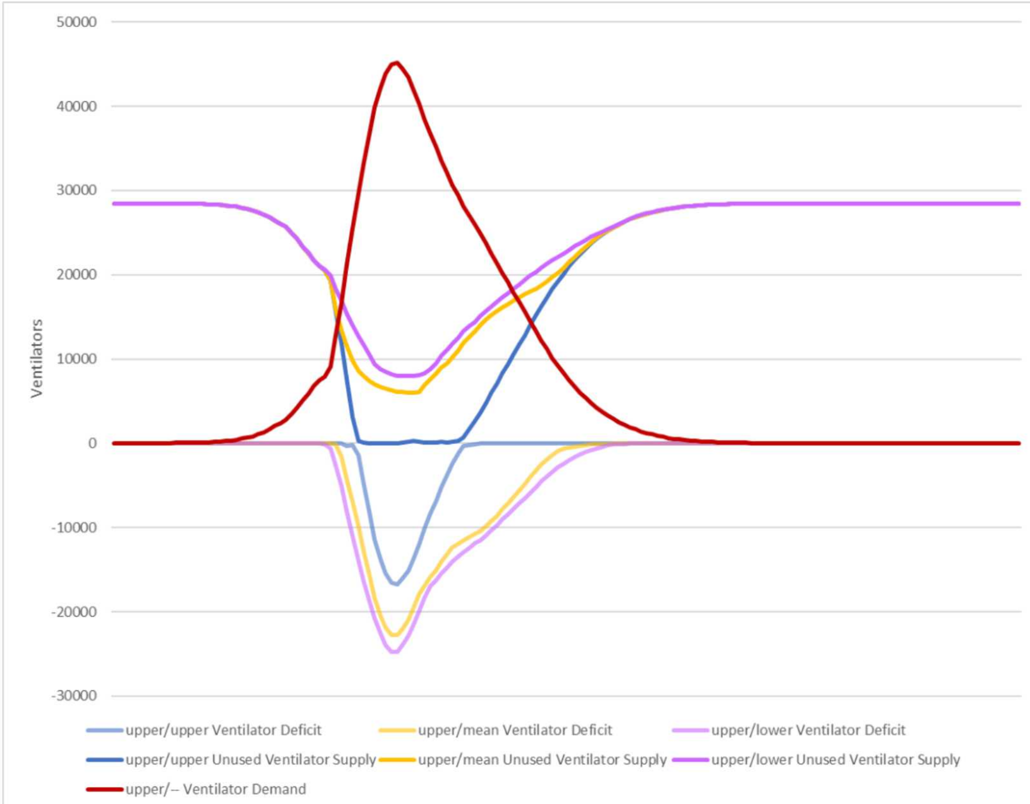


**Figure 6-2. Ventilator demand, deficit, and unused supply time-series for the scenarios using IHME Upper patient streams for observations.**

While care must be taken not to make general rules from such a focused analysis, these results indicate the potential for this framework to provide useful information to policymakers. In this case, it is clear that there appears to be significantly worse outcomes if one uses "optimistic" forecasts when determining if nation-wide ventilator routing is feasible. On the other hand, while using "worst-case" forecasts does not always results in near optimal performance, they do perform fairly well, especially considering they are using forecasts which do not match what actually happens. From a policy perspective, this indicates that when trying to make decisions for an uncertain future about the feasibility and potential implementation for regional ventilator routing, choosing the "worst-case" (IHME Upper) patient projections (which perform the best against all observed patient scenarios) is probably the best course of action.

## 6.2. Regionality in Ventilator Movement Routes

Another output from the ISDR ventilator model is the resource movement routes specified by the routing model. Drawing too much insight from specific movements should be avoided as the model's relative simplicity is focused on determining feasibility and not prescribing specific logistic plans. However, aggregating the movements over an entire model run can provide useful information about how ventilators are being shared across the country. This information may assist policymakers in understanding which states may benefit by building partnerships for sharing resources. An example allocation solution aggregated across the entire simulation time is shown in

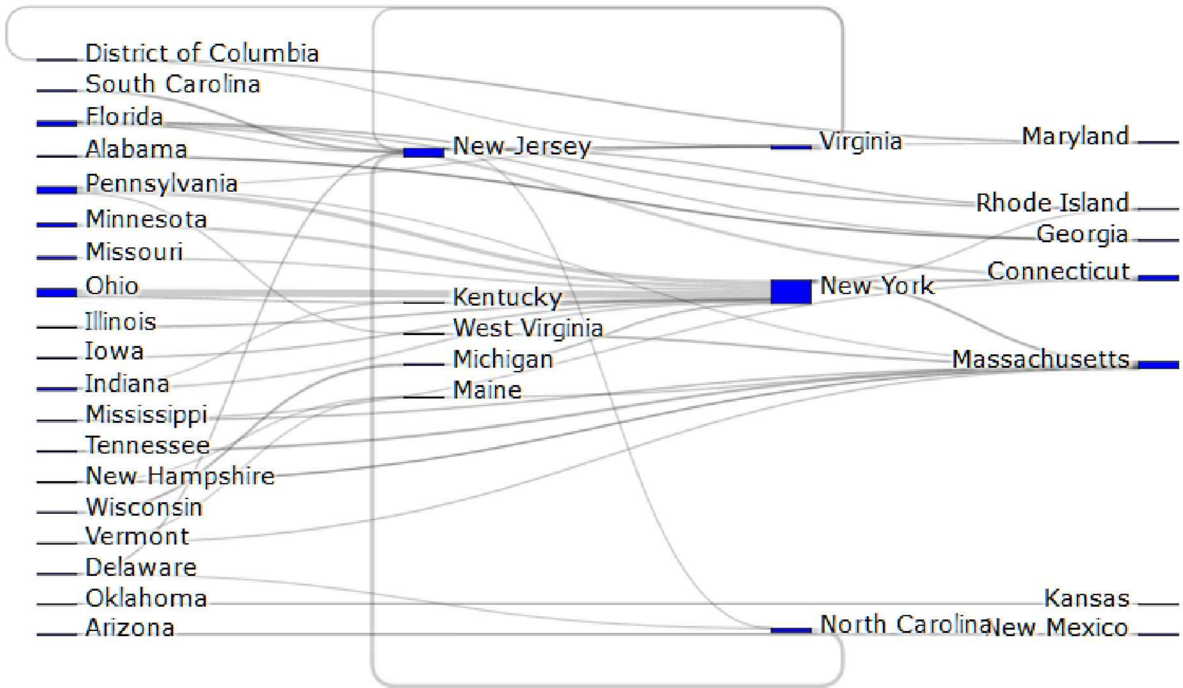


**Figure 6-3. Ventilator routing recommendations across the entire simulation time to minimize the shortfall experienced by any state. This Sankey diagram connects nodes that share resources.**

Figure 6-3, which connects states by the flows between them. It is important to note the marked lack of cyclic structures in the figure. This indicates the solution determined by the routing algorithm avoided scenarios where one state donated ventilators, only to need them back in the future. This sort of 'give up what you have now, even though you may need it back later' mentality is one that local government officials may be too risk averse to have, given the potential fall out of giving away lifesaving resources only to need them later. Time aggregate analyses like this can bolster confidence in collaborative agreements and be used to apply attribution toward a strategic decision in cooperative contexts.

However, these routing recommendations are temporally explicit, and over a given time period, the relationships between pairs or groups of states will self-organize to mitigate shortfalls while minimizing the costs associated with shipments. These shipments construed as a weighted directed network can be viewed at discrete points in time to illustrate the kinds of dynamics required to achieve the shortfall mitigation performance described in Table 6-2. Figure 6-4 illustrates the routing recommendations for a single time point at the national scale.

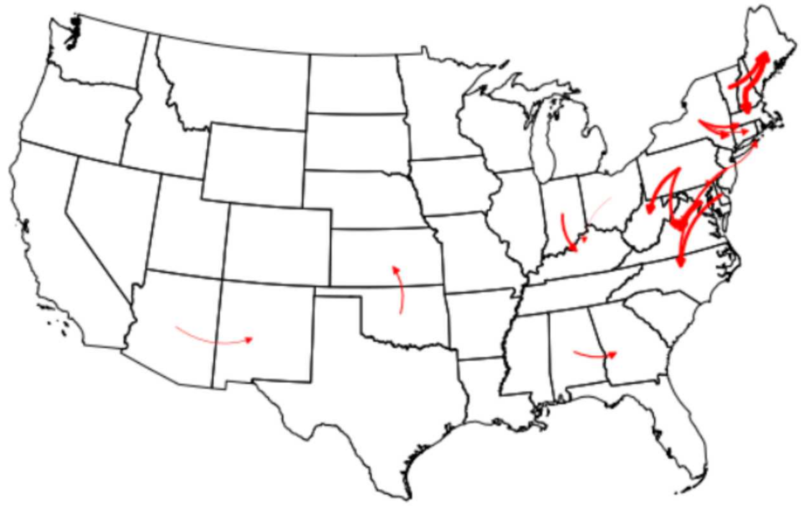

**Figure 6-4. Ventilator routing recommendations to minimize the shortfall experienced by any state for a single time point. Arrows represent the direction of resource flow, weighted by the magnitude of the shipment.**

The time aggregated product of this model output then can be generated by summing ventilator movements across all states and times. The summation ignores the directionality of the movements since this does not affect potential interstate partnerships. This relationship can then be construed as an undirected network, with locations representing nodes and the resources shared between them as edges, the adjacency matrix of which is weighted by the magnitude of the commodity transfer. We sought to leverage this network structure to determine which states tended to cooperate by effectively clustering their transactions. This cluster (or community) detection then represents the logical partnerships – in the context of ventilator sharing – that states may want to make.

We chose the Louvain community clustering algorithm to determine these partnerships. This algorithm optimizes a modularity score that effectively seeks to increase the strength of intra-community connections and reduce the strength of inter-community connections [4]. Figure 6-3 through Figure 6-7 show the result of running the Louvain algorithm for each of the model scenarios. For the runs using IMHE Lower and Mean patient flows for the observations, the same communities were identified for all three forecasts. For the runs using IMHE Upper streams for the observations, three different community maps were generated, though they all share some general similarities with one another.

There is a fair amount of variation in community membership across the model runs, suggesting the partnerships recommended by the routing algorithm are heavily influenced by the patient and demand estimates. However, several themes emerge despite the variability inherent scenario differences. Generally, communities tend to be relatively spatially compact, which points towards the idea that partnerships of neighboring states may make the most sense. While this is likely a function of the cost function equating to distance, consortiums of states and organizational units (e.g., FEMA regions) tend to be geospatially clustered as well.

Additionally, the south-eastern states and north-eastern states extending to Ohio tend to group into a limited number of communities. This suggests that each of these regions may individually benefit from forming interstate partnerships, and, given the mixing at the boundaries between them (e.g.

Massachusetts and Virginia belong to both in different scenarios) perhaps an entire eastern U.S. partnership would make sense. For the mid-west and western states, the picture is a little less clear, but again there is a tendency for the states to build communities with neighboring states.

The broader point is that the results of this modeling analysis is that regional partnerships should at least be considered, and results such as those shown in these maps could help policymakers make strategic decisions in anticipation of needs in future crises. Further analysis could include trying other community detection algorithms, as well as quantifying the relative strength of the identified communities to better assess their potential for forming partnerships.

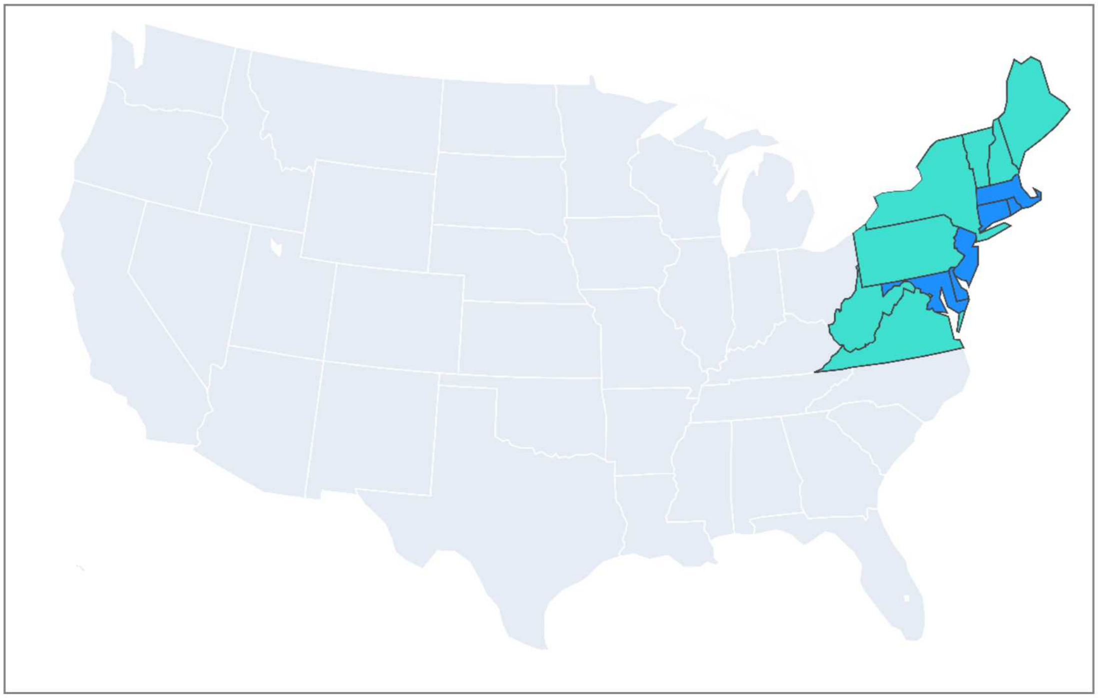

**Figure 6-5. Louvain communities detected for ventilator movements between states for the scenarios using IMHE Lower patient flows for observations. States that belong to the same community share the same color.**

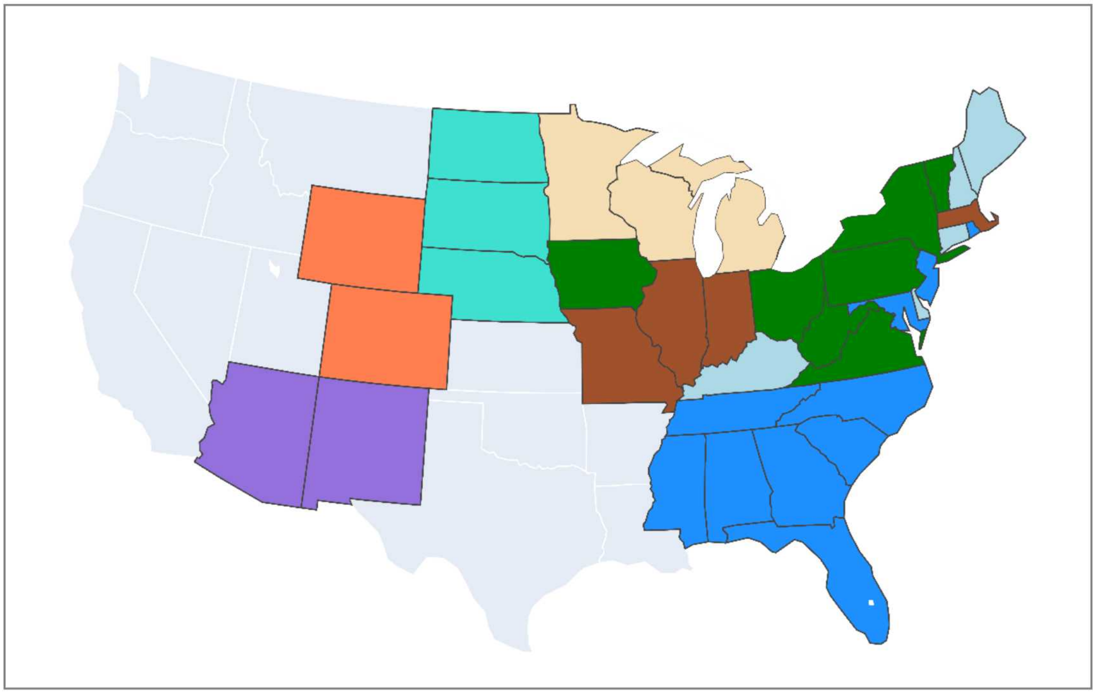

**Figure 6-6. Louvain communities detected for ventilator movements between states for the scenarios using IMHE Mean patient flows for observations. States that belong to the same community share the same color.**

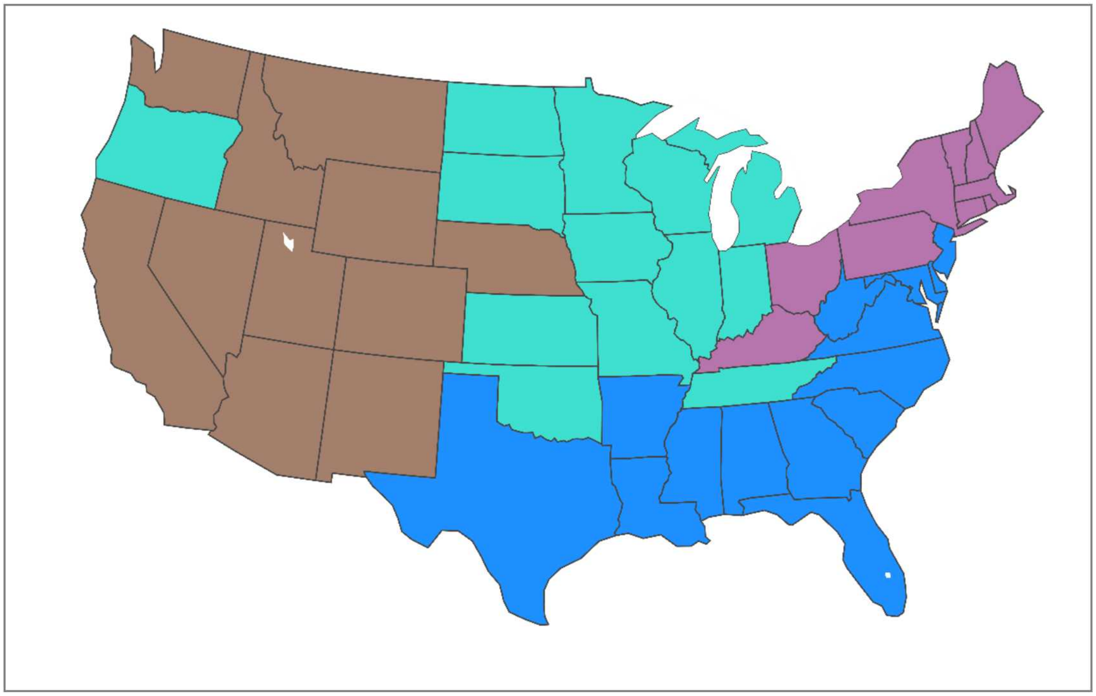

**Figure 6-7. Louvain communities detected for ventilator movements between states for the upper/lower scenario. States that belong to the same community share the same color.**

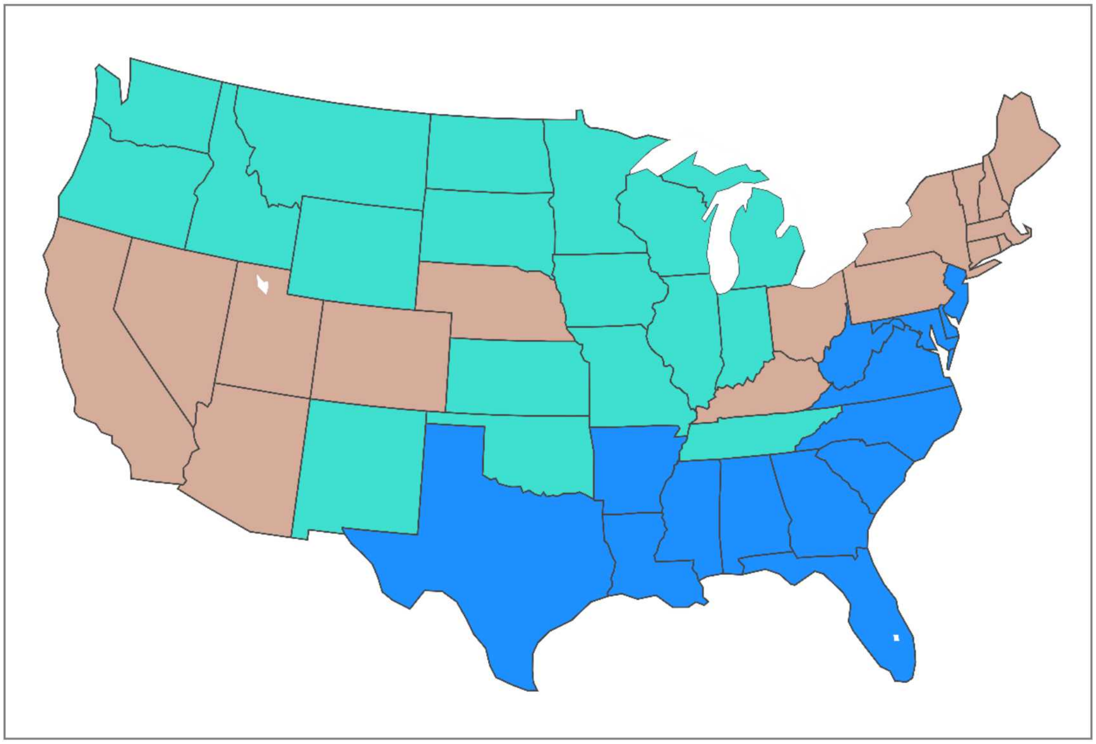

**Figure 6-8. Louvain communities detected for ventilator movements between states for the upper/mean scenario. States that belong to the same community share the same color.**

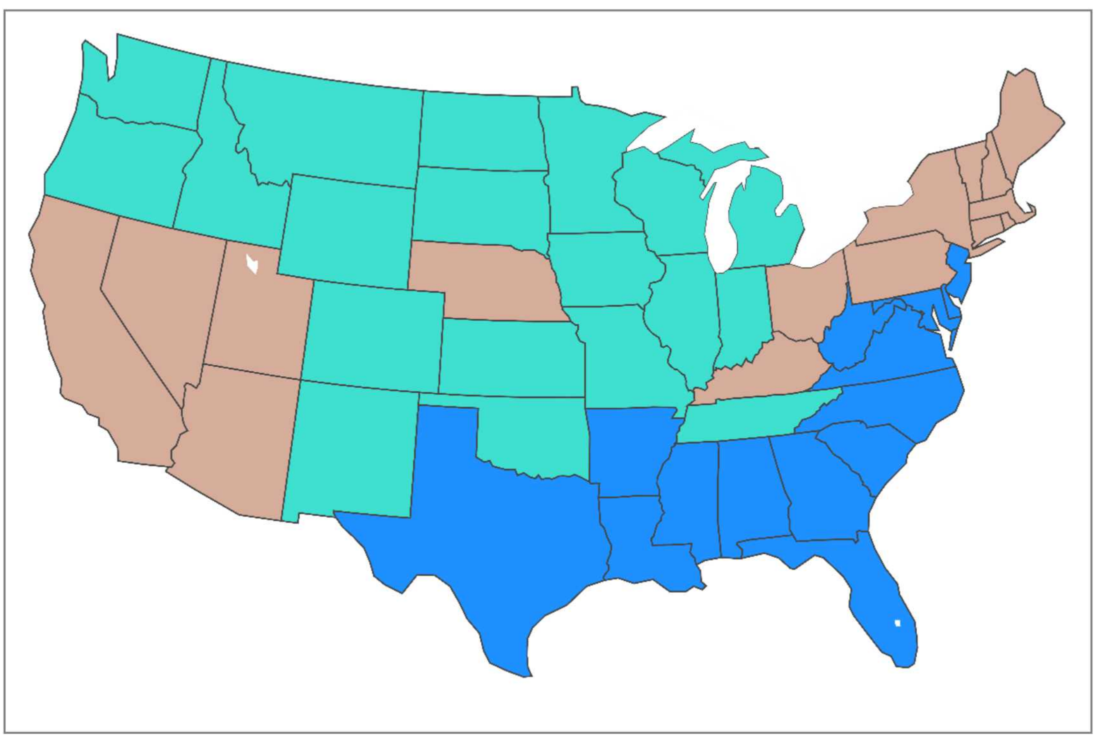

**Figure 6-9. Louvain communities detected for ventilator movements between states for the upper/upper scenario. States that belong to the same community share the same color.**

# 7. EFFECTS OF UNCERTAINTY

As mentioned in section 3, the treatment demand model introduces uncertainty in a number of places as patients are processed. In the previous section, the model runs all used the same random seed in order to focus the analysis on comparability. This section will look at varying the random seed so as to more fully understand the effects of uncertainty on the results. It will become clear that the variation introduced by the uncertainty is relatively small compared with the results of interest, and thus using a single, fixed-seed run for analysis is a valid simplification.

To measure the uncertainty effects, the ISDR ventilator model was run 100 times, each using a different random seed, for the scenario using the IMHE upper patient streams for the observations and the IMHE mean streams for the forecasts. For each model day, the minimum, maximum, and average values of the results across all the runs is calculated. Graphs of these results for the ventilator demand, ventilator deficit, and unused ventilator supply are shown in Figure 7-1, Figure 7-2, and Figure 7-3. Ventilator demand variation is the direct effect of the treatment model uncertainty, and the ventilator deficit and unused supply are indirect results from the supply routing model. As seen in the charts, there is not a large variation in the results due to model uncertainty, as both the shape of the curves and their comparative values closely track one another.

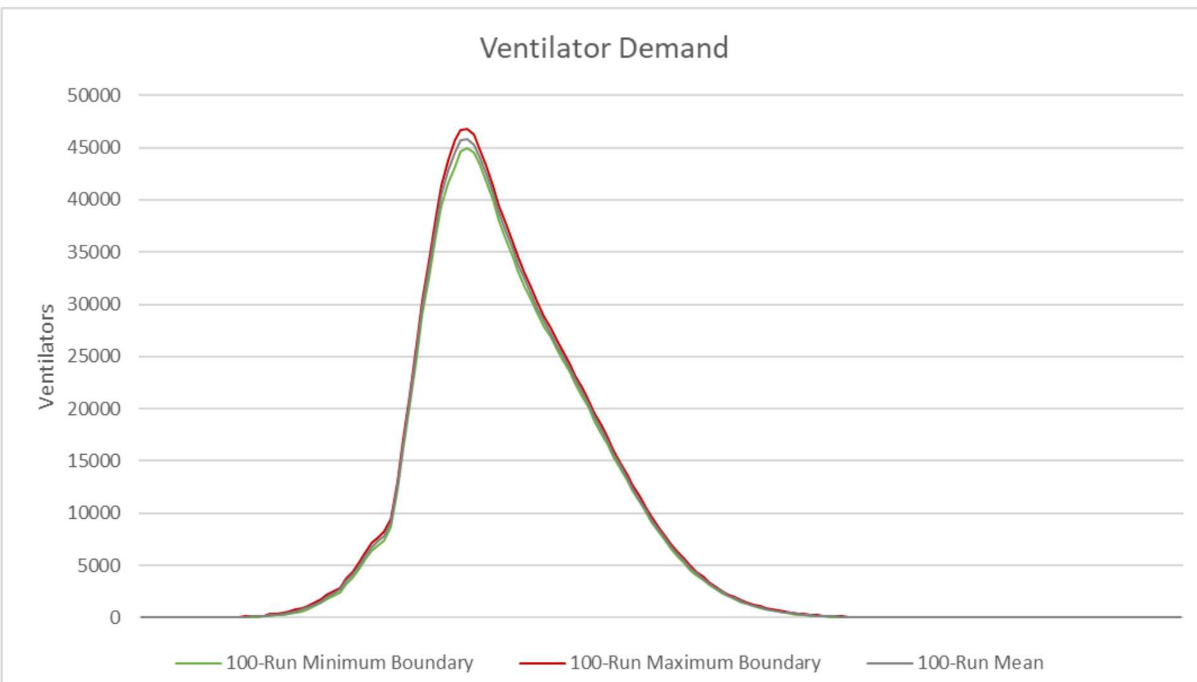


**Figure 7-1. Range and mean of daily ventilator demand for 100-run uncertainty analysis using IMHE upper for observations and IMHE mean for forecasts.**

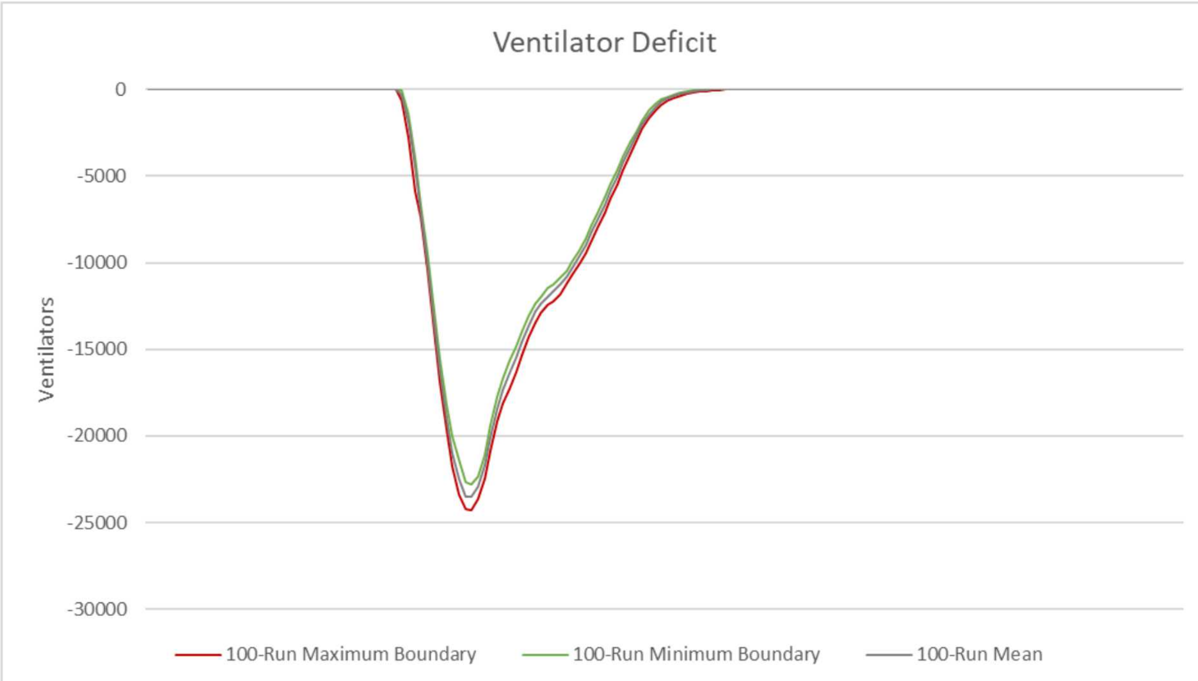


**Figure 7-2. Range and mean of daily ventilator deficits for 100-run uncertainty analysis using IMHE upper for observations and IMHE mean for forecasts.**

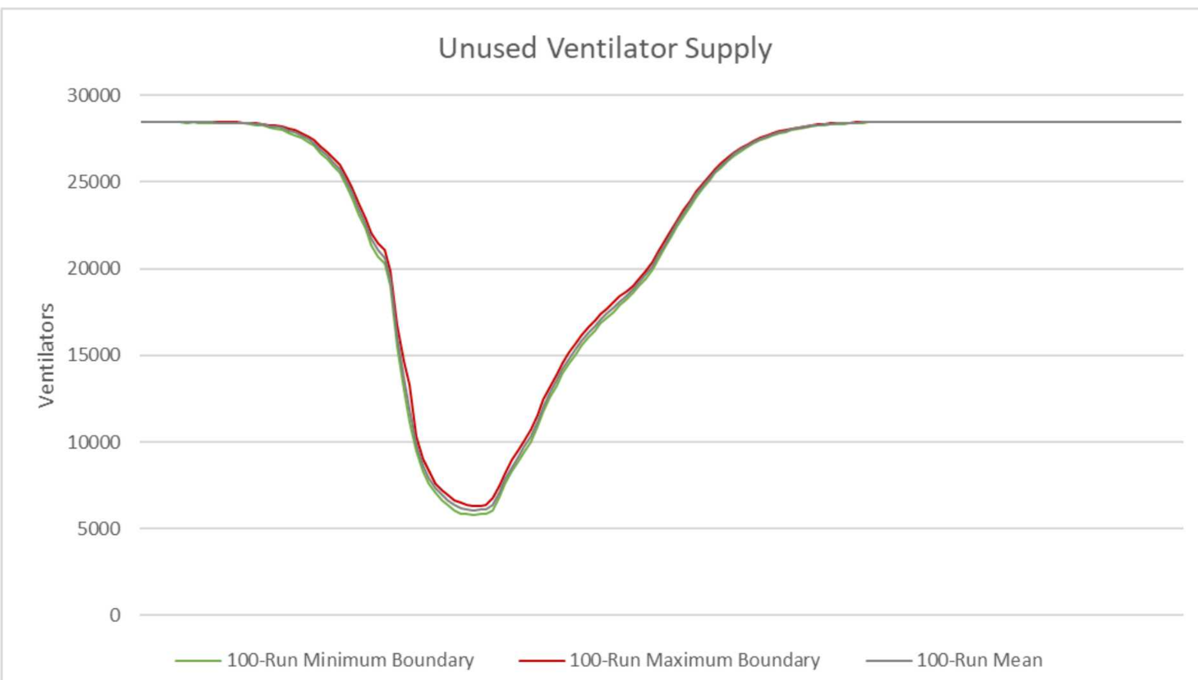


**Figure 7-3. Range and mean of daily unused ventilator supply for 100-run uncertainty analysis using IMHE upper for observations and IMHE mean for forecasts.**

Table 7-1 quantifies these comparisons, showing the greatest spread between the maximum and minimum boundaries as an absolute value and a percentage of the mean values - both the mean value at the day of greatest spread and the largest mean value in the chart. While the spreads are not insignificant, they are reasonably small enough to not invalidate the conclusions of the previous section.

**Table 7-1. Summary of extreme variations in results for 100-run uncertainty analysis using IMHE upper for observations and IMHE mean for forecasts.**

| Result Metric | Greatest spread between maximum and minimum boundary | Spread value as percentage of mean value at that time | Spread value as percentage of maximum mean value |
|---|---|---|---|
| Ventilator Demand | 2521 | 5.7% | 5.5% |
| Ventilator Deficit | -2018 | 9.0% | 8.6% |
| Unused Ventilator Supply | 2082 | 17.6% | 7.3% |

As a specific comparison, Figure 7-4 shows the previous 100-run chart for unused ventilator supply with the results from the previous section's analysis for the IMHE upper observation runs. The mean forecast run closely tracks the 100-run mean, whereas the lower and upper forecast runs both clearly lie outside of the 100-run minimum and maximum boundaries.

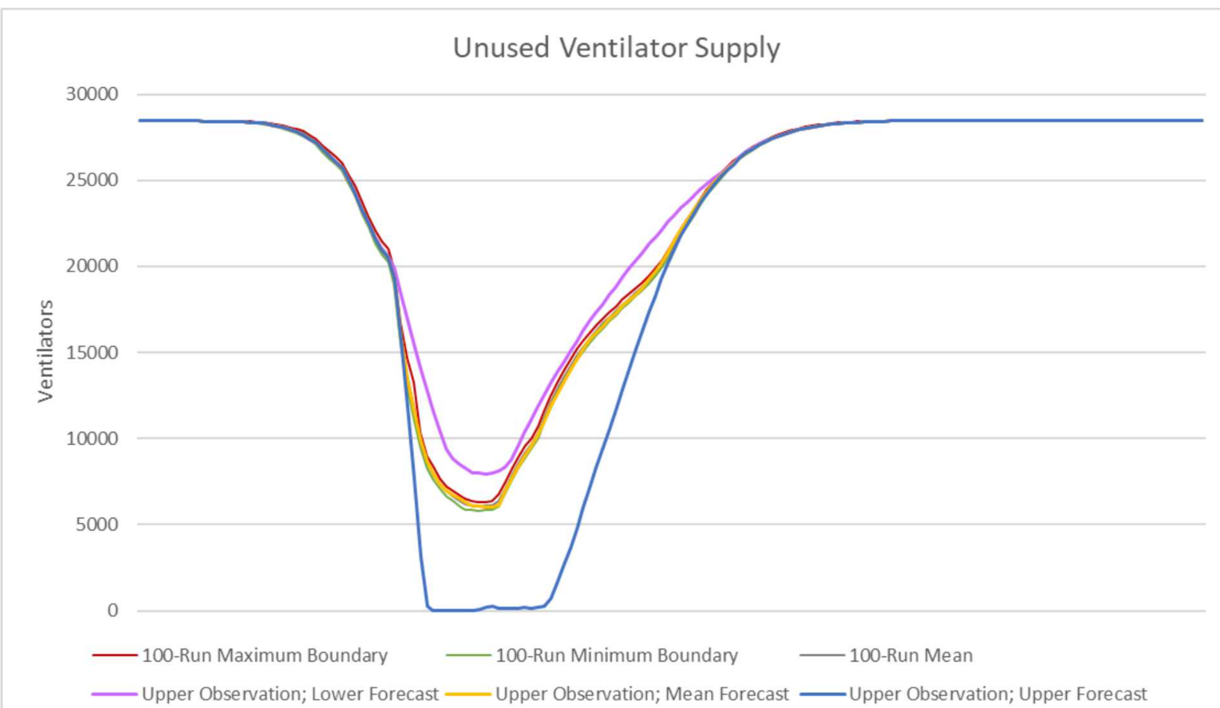


**Figure 7-4. Daily unused ventilator supply results for 100-run uncertainty analysis using IMHE upper for observations and IMHE men for forecasts, and specific model run results for scenarios using IMHE upper for observations.**

# 8. CONCLUSION

This report discusses the details of the ISDR framework along with two models adapted for use within it: a treatment demand model and an optimizing supply routing model. An application of the integrated model to ventilators for COVID-19 patients demonstrates the capability of the models to provide useful policy insights to stakeholders. The insights include assessing the feasibility of sharing ventilators between states, better understanding the implications of choosing a specific forecast for making policy decisions, and recognizing the interstate partnerships that may help to mitigate future resource crises.

While this capability could be simply adapted for use with other resources, there are also a number of directions that this work could be extended for future applications. One is to incorporate more resources at once within ISDR; doing so would not only capture a fuller picture of healthcare resource supply and demand, but also more complex models and routers could be built which could balance tradeoffs – such as economics and resilience – to help policymakers better understand the decision space.

Another extension would be to incorporate meaningful uncertainty quantification into the model, particularly around the parameter space of the input models. As evidenced by the COVID-19 epidemic, during a crisis it can be difficult to ascertain even reasonable parameter estimates for models, and being able to quickly and usefully assess the consequences of this uncertainty is quite important.

A final extension involves the use of advanced visualization and data exploration techniques to allow analysts and stakeholders a more direct way to interact with the model outputs. The technical knowledge needed to set up and run these models can be a barrier that extends into the analysis of the results. By providing better accessibility and tooling, these unnecessary technical barriers can be removed and allow talented analysts and important policymakers to assist in the development of knowledge and insights from the models.

# REFERENCES


[1] L. P. Swiler, T. Portone and W. E. Beyeler, "Uncertainty Analysis of Hospital Resource Demand Model for Covid-19," SAND2020-4900. Sandia National Laboratories, Albuquerque, NM, May, 2020.

[2] New York State Governor's Office, "Governor Cuomo is a Guest on MSNBC's The Rachel Maddow Show," 2 April 2020. [Online]. Available: https://www.governor.ny.gov/news/audio-rush-transcript-governor-cuomo-guest-msnbcs-rachel-maddow-show. [Accessed 10 May 2020].

[3] Institute for Health Metrics and Evaluation, "Covid-19 Estimate Downloads," [Online]. Available: http://www.healthdata.org/covid/data-downloads. [Accessed 7 April 2020].

[4] V. D. Blondel, J.-L. Guillaume, R. Lambiotte and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of Statistical Mechanics: Theory and Experiment,* vol. 2008, no. 10, p. P10008, 2008.

# DISTRIBUTION

**Email—Internal**

| Name | Org. | Sandia Email Address |
|---|---|---|
| Laura P. Swiler | 1463 | lpswile@sandia.gov |
| Walter E. Beyeler | 6617 | webeyel@sandia.gov |
| Katherine D. Tremba | 6824 | kdtremb@sandia.gov |
| Dean A. Jones | 8720 | dajones@sandia.gov |
| Sean E. DeRosa | 8721 | sderosa@sandia.gov |
| Patrick D. Finley | 8721 | pdfinle@sandia.gov |
| Christopher R. Frazier | 8721 | crfrazi@sandia.gov |
| Jared L. Gearhart | 8721 | jlgearh@sandia.gov |
| Stephen D. Kleban | 8722 | sdkleba@sandia.gov |
| Daniel J. Krofcheck | 8722 | djkrofc@sandia.gov |
| Katherine A. Klise | 8825 | kaklise@sandia.gov |
| Technical Library | 01977 | sanddocs@sandia.gov |

This page left blank

This page left blank

Sandia National Laboratories
is a multimission laboratory
managed and operated by
National Technology &
Engineering Solutions of
Sandia LLC, a wholly owned
subsidiary of Honeywell
International Inc. for the U.S.
Department of Energy's
National Nuclear Security
Administration under contract
DE-NA0003525.