

## ANESTHESIOLOGY

# Risk Stratification Index 3.0, a Broad Set of Models for Predicting Adverse Events during and after Hospital Admission

Scott Greenwald, Ph.D., George F. Chamoun, B.S., Nassib G. Chamoun, M.S., David Clain, B.S., Zhenyu Hong, M.S., Richard Jordan, Ph.D., Paul J. Manberg, Ph.D., Kamal Maheshwari M.D., Daniel I. Sessler, M.D.

*ANESTHESIOLOGY* 2022; 137:673–86

## EDITOR'S PERSPECTIVE

### What We Already Know about This Topic

- Uses for risk stratification tools include setting baselines for health service evaluations, identifying patients who may need higher levels of care, and allocating hospital resources

### What This Article Tells Us That Is New

- From a dataset of more than 9 million patients, a risk score based on administrative claims history was developed to provide individualized risk profiles at hospital admission that may help guide patient management

## ABSTRACT

**Background:** Risk stratification helps guide appropriate clinical care. Our goal was to develop and validate a broad suite of predictive tools based on International Classification of Diseases, Tenth Revision, diagnostic and procedural codes for predicting adverse events and care utilization outcomes for hospitalized patients.

**Methods:** Endpoints included unplanned hospital admissions, discharge status, excess length of stay, in-hospital and 90-day mortality, acute kidney injury, sepsis, pneumonia, respiratory failure, and a composite of major cardiac complications. Patient demographic and coding history in the year before admission provided features used to predict utilization and adverse events through 90 days after admission. Models were trained and refined on 2017 to 2018 Medicare admissions data using an 80 to 20 learn to test split sample. Models were then prospectively tested on 2019 out-of-sample Medicare admissions. Predictions based on logistic regression were compared with those from five commonly used machine learning methods using a limited dataset.

**Results:** The 2017 to 2018 development set included 9,085,968 patients who had 18,899,224 inpatient admissions, and there were 5,336,265 patients who had 9,205,835 inpatient admissions in the 2019 validation dataset. Model performance on the validation set had an average area under the curve of 0.76 (range, 0.70 to 0.82). Model calibration was strong with an average  $R^2$  for the 99% of patients at lowest risk of 1.00. Excess length of stay had a root-mean-square error of 0.19 and  $R^2$  of 0.99. The mean sensitivity for the highest 5% risk population was 19.2% (range, 11.6 to 30.1); for positive predictive value, it was 37.2% (14.6 to 87.7); and for lift (enrichment ratio), it was 3.8 (2.3 to 6.1). Predictive accuracies from regression and machine learning techniques were generally similar.

**Conclusions:** Predictive analytical modeling based on administrative claims history can provide individualized risk profiles at hospital admission that may help guide patient management. Similar results from six different modeling approaches suggest that we have identified both the value and ceiling for predictive information derived from medical claims history.

(*ANESTHESIOLOGY* 2022; 137:673–86)

This article is featured in "This Month in Anesthesiology," page A1. This article is accompanied by an editorial on p. 661. This article has a related Infographic on p. A17. Supplemental Digital Content is available for this article. Direct URL citations appear in the printed text and are available in both the HTML and PDF versions of this article. Links to the digital files are provided in the HTML text of this article on the Journal's Web site ([www.anesthesiology.org](http://www.anesthesiology.org)). This article has an audio podcast. This article has a visual abstract available in the online version. The work presented in this article has been presented at the American Society of Anesthesiologists Annual Meeting in New Orleans, Louisiana, October 22, 2022.

Submitted for publication March 24, 2022. Accepted for publication August 31, 2022. Published online first on September 21, 2022.

Scott Greenwald, Ph.D.: Health Data Analytics Institute, Dedham, Massachusetts.

George F. Chamoun, B.S.: Health Data Analytics Institute, Dedham, Massachusetts.

Nassib G. Chamoun, M.S.: Health Data Analytics Institute, Dedham, Massachusetts.

David Clain, B.S.: Health Data Analytics Institute, Dedham, Massachusetts.

Zhenyu Hong, M.S.: Health Data Analytics Institute, Dedham, Massachusetts.

Richard Jordan, Ph.D.: Health Data Analytics Institute, Dedham, Massachusetts.

Paul J. Manberg, Ph.D.: Health Data Analytics Institute, Dedham, Massachusetts.

Kamal Maheshwari, M.D.: Department of Outcomes Research, and Department of General Anesthesiology, Cleveland Clinic, Cleveland, Ohio.

Daniel I. Sessler, M.D.: Department of Outcomes Research, Cleveland Clinic, Cleveland, Ohio.

Copyright © 2022, the American Society of Anesthesiologists. All Rights Reserved. *Anesthesiology* 2022; 137:673–86. DOI: 10.1097/ALN.0000000000004380

Risk stratification tools are useful in at least four distinct situations. The first is health services research. Specifically, comparisons among various facilities or treatment groups can only fairly be evaluated after adjusting for baseline risk of the outcomes of interest. For example, health services evaluations comparing mortality among various hospitals must adjust for baseline mortality risk and procedural complexity across the relevant populations.<sup>1</sup> Accurate risk stratification similarly contributes to observational research by providing an accurate basis for propensity matching and multivariable regression. The second role for risk stratification is to identify enriched populations for care pathways and clinical trials<sup>2</sup>—that is, selecting patients most likely to experience an adverse event and benefit from specific interventions. The third role for risk stratification is to guide clinical care, including decisions about which surgical or alternative treatment options are most likely to prove helpful.<sup>3</sup> Finally, reliable predictions of expected duration of hospitalization and discharge disposition can help guide hospital resource management and planning for follow-up support services.<sup>4</sup>

The Risk Stratification Index version 1.0, first introduced in 2010, was a broadly applicable risk adjustment measure for predicting mortality in-hospital, at 1 month, and at 1 yr.<sup>5</sup> The original models were derived from more than 35 million Medicare hospitalizations between 2001 and 2006, and were thereafter validated in a wider age range of California inpatients<sup>6</sup> and in two single-center studies.<sup>5,7,8</sup> The index also performed well in an independent set of 39 million Medicare admissions from 2008 to 2012.<sup>9</sup> Version 2.0 of the Risk Stratification Index,<sup>10</sup> introduced in 2018, used the expanded set of International Classification of Diseases, Ninth Revision, codes and information that Medicare now allows for each admission, including up to 25 diagnostic codes, 25 procedure codes, and flags indicating conditions that were present on admission. Both Risk Stratification Index versions provided better discrimination than the Charlson Comorbidity Index and other publicly available stratification systems based on administrative data.<sup>5,10</sup> Version 2.0 was also well calibrated.<sup>9</sup> While these versions have been used for academic research,<sup>11–16</sup> we are not aware that they are being used for clinical care.

Previous versions of the Risk Stratification Index were based exclusively on diagnosis and procedure codes from the index hospitalization, and relied heavily on present-on-admission codes. The difficulties with this approach are that billing codes and present-on-admission flags are usually generated by specialized coders after patients are discharged. Consequently, key information necessary for accurate risk stratification is generally unavailable at the time of hospital admission—when stratification may be especially useful. An additional consequence of basing stratification on a single hospitalization is that temporally restricted information fails to capture individuals' preadmission illness trajectories, which might improve predictions. Another limitation

of previous versions of the Risk Stratification Index is that they were based on International Classification of Diseases, Ninth Revision, codes, rather than International Classification of Diseases, Tenth Revision, codes, which are now universally used. Previous versions were also restricted to surgical admissions rather than also considering medical admissions. Finally, previous Risk Stratification Index models were restricted to in-hospital, 30-day, and 1-yr mortality, along with hospital length of stay.

Our primary goal was therefore to develop and validate a broad suite of practical analytic tools based on International Classification of Diseases, Tenth Revision, diagnostic and procedural code histories for predicting hospital utilization outcomes and adverse events for both surgical and medical inpatient admissions. Specifically, we derived predictors for meaningful utilization endpoints including unplanned hospital admissions, discharge status, and excess length of stay, along with major adverse events and complications including in-hospital and 90-day mortality, acute kidney injury, sepsis, pneumonia, respiratory failure, and a composite of major cardiac complications.

As in previous versions of the Risk Stratification Index, we primarily used logistic regression because the method provides easily interpretable results including model coefficients that identify key drivers of risk and quantify their relative contributions. However, machine learning methods have become increasingly popular and have shown better predictive perioperative performance than clinical scoring systems in some but not all settings.<sup>17–19</sup> We therefore developed analogous models using five commonly used machine learning methods and compared model performance characteristics with each method.

## Materials and Methods

Our primary analyses were conducted on the Centers for Medicare and Medicaid Services (Baltimore, Maryland) Research Identifiable File data on a remote server under a Centers for Medicare and Medicaid Services data use agreement (No. 51870). Access to the remote server was provided through VM Horizon Client (5.3; VMware, Inc, USA) and analysis conducted using SAS (9.04; SAS Institute, USA) within SAS Enterprise Guide (7.15; SAS Institute). Secondary (robustness) analyses were conducted on a 5% sample (Limited Data Set) of the same Medicare dataset provided by the Centers for Medicare and Medicaid Services housed on a local server using R software (version 4.2.0; available at <https://cran.r-project.org/src/base/release/2022-04-22>) under a separate data use agreement (LDSS-2017-51396). Data were handled consistent with our data use agreements, which required suppression of metrics in downloaded tables for populations smaller than 11 individuals.

Our analyses were determined to be exempt from informed consent requirements by the New England Institutional Review Board (Needham, Massachusetts). This

report follows the Transparent Reporting of a Multivariable Prediction Model for Individual Prognosis or Diagnosis reporting guideline.<sup>20</sup>

The bulk of our statistical analysis plan was submitted to the Centers for Medicare and Medicaid Services in response to their Artificial Intelligence Health Outcome Challenge in February 2020, before formal analysis began. The plan included primary use of logistic regression, specific endpoints, the metrics to be reported, reporting results at outcome incidence and at the top 5%, and *a priori* definitions of adequate model performance. A comparison to various machine learning methods was anticipated, although the specific methods were not prespecified. Lift (enhancement ratio) was not part of the original plan, and was added during the analysis, which was conducted from October 2021 to January 2022.

### Subject Selection

We used the full Medicare fee-for-service and dual-eligible (Medicaid and Medicare) files for beneficiaries hospitalized in 2017 to 2019. Admissions were excluded if patient age on admission was younger than 18 or older than 99 yr, records had missing or inconsistent data (e.g., missing sex or birthdate information, or had different sex, birth dates, or mortality dates [if applicable] reported in source files), or patients had either discontinuous Part A or Part B Medicare coverage or had Part C coverage in the year before admission. Claims data during the year before the admission were used to characterize the patient history. Claims data during the 90 days after admission characterized outcomes. The admission status was classified as “planned” if the reason for admission was elective, and otherwise designated “unplanned.”

### Outcomes Selection

We present a suite of 10 models that predict excess length of stay and adverse events, selected to demonstrate performance of predictors for clinically and economically meaningful outcomes spanning a broad range of incidences. Cardiac complications, kidney injury, sepsis, pneumonia, and respiratory failure were defined using International Classification of Diseases, Tenth Revision, diagnosis and procedure codes<sup>21</sup> along with information about their associated claim, such as the setting and revenue center. Additionally, we considered whether codes were primary or secondary.

Endpoint definitions were derived using published methods for classifying events using administrative data. Event definitions and their associated references are presented in Supplemental Table 1 (<http://links.lww.com/ALN/C923>). Events were identified between admission and discharge (for in-hospital endpoints) and/or between admission and 90 days thereafter (for 90-day endpoints). A 90-day observation window was chosen for events and

mortality because previous reports suggest that 90-day outcomes may be more reliable than 30-day outcomes for measuring hospital performance.<sup>22–24</sup> In-hospital mortality was defined by any-cause death between admission and discharge. Ninety-day mortality was defined by death between admission and 90 days thereafter. Excess length of stay was defined as the difference between the observed duration of hospitalization and the geometric mean length of stay associated with the default 2021 v1 Clinical Classifications Software Refined<sup>25</sup> category for the admitting diagnosis when the admission was unplanned, or the 2020 Clinical Classifications Software<sup>26</sup> category associated with the primary procedure for planned admissions. Discharge status to a facility was defined by discharge to locations other than home, with or without organized home health care.<sup>27</sup>

### Model Development

Medical history was represented by a set of variables indicating the presence or absence of individual and categories of International Classification of Diseases, Tenth Revision, diagnostic and procedure codes. We used a custom procedure to reduce 69,000 potential International Classification of Diseases, Tenth Revision, diagnostic codes to a representative subset of 4,426 codes by collapsing rare codes into their parent codes to avoid overfitting (Supplemental Figure 1, <http://links.lww.com/ALN/C923>). International Classification of Diseases, Tenth Revision, diagnostic codes were additionally represented by their corresponding default Clinical Classifications Software Refined category.<sup>25</sup> Similarly, International Classification of Diseases, Tenth Revision, procedure codes were represented by their corresponding default Clinical Classifications Software category.<sup>26</sup> Temporal information relative to a prediction date was encoded using two sets of these variables representing the presence or absence of relevant codes in the past 90 or 365 days.

Outcomes were indexed to the date of inpatient admission, and claims within the preceding 365 days were included in our models. The only information used from the day of admission was the admitting diagnosis (rather than the principal diagnosis) along with the principal procedure for planned admissions. We also included age at the time of admission.

Logistic regression models were trained with the SAS HPLOGISTIC procedure using log-log linkage and backwards fast selection of covariates, keeping those with a  $P < 0.01$  significance level. We used the asymmetric log-log link function because such models handle skewed extreme value distributions associated with rare events better than symmetrical link functions.<sup>28</sup> There were nonlinear interactions by sex and admission type between International Classification of Diseases, Tenth Revision, codes and various outcomes which preclude using a single logistic model for each outcome. We therefore constructed an overall

process to apply the appropriate model coefficients from the ensemble of four models depending on sex and admission status.

### Model Application

Our general approach was to (1) train a model on 80% of qualifying Medicare admissions from 2017 and 2018 from the development dataset (training set); (2) apply the resulting model to the remaining 20% of the development dataset (test set) to document modeling robustness; and (3) prospectively evaluate the resulting final model on out-of-sample admissions from 2019 to document model validation (prospective validation set).

### Performance Metrics

Overall discrimination performance was evaluated using the area under the receiver operating characteristics curves (AUC). Performance at a given operating threshold was assessed by sensitivity, positive predictive value, and lift.<sup>29</sup> Lift is defined as the ratio of detected events using the classifier relative to not using the classifier, which is equivalent to the positive predictive value divided by the incidence. When predictive classifiers are used to identify an enriched subpopulation, lift therefore quantifies the enrichment ratio.

To compare model detection performance consistently across various endpoints, we compared sensitivity, positive predictive value, and lift for each model at an alert threshold corresponding to the highest 5% risk fraction of the population. This sort of high-risk threshold might be used clinically to identify subpopulations most likely to benefit from intervention. We similarly compared sensitivity, positive predictive value, and lift for each model at thresholds corresponding to the observed incidence for each endpoint within the population. Evaluating detector performance using endpoint-specific thresholds normalizes performance results, thereby simplifying performance comparisons across various endpoints and modeling methods. We did not try to identify thresholds that optimize positive and negative predictive values because optimizing depends upon the relative costs of false detections and missed events, which are specific to individual endpoints and use cases. Specifically, selection of the most appropriate thresholds represents a form of resource constrained ranking,<sup>12</sup> where selection of a particular threshold to define a “higher-risk” subgroup from a ranked population is based on a particular use case.

One hundred bins in steps of 1% resolution of risk were used to identify subpopulations along the full continuum of risk. To evaluate calibration, we computed correlation coefficients between observed and predicted incidences for each endpoint using all bins having more than 100 subjects. We similarly computed observed-to-predicted ratios. For the excess length of stay model, the only nonbinary event in the suite, we estimated the root mean squared error of the absolute predictions *versus* the observed values in groups

of 1,000 individuals. We computed the mean and 95% CI for AUC.

We *a priori* set conservative minimum acceptable performance criteria using two metrics to reject clinically non-viable models. Model acceptance required (1) a reasonably accurate overall classification performance defined by an AUC 0.70 or greater and (2) relatively accurate prediction defined by an observed-to-expected ratio near 1 over the full risk continuum (*i.e.*, calibration  $R^2$  greater than 0.80). The conservative 0.7 minimum acceptance threshold for AUC was based on consultation with clinical advisors and a literature review indicating the acceptability of numerous perioperative machine learning models with *c*-statistics in the 0.7 to 0.8 range.<sup>18,30</sup> Because no *a priori* hypotheses were tested, we did not estimate required sample size, and instead used all eligible cases available in the Medicare fee-for-service files for the selected years. To evaluate the importance of endpoint-specific models, we compared incidence of various complications in patients selected for having the highest 5% risk of 90-day mortality to those with the highest 5% risk of specific complications.

### Model Comparisons

In addition to our primary models, which were developed using multivariable regression, we developed models based on five machine learning methods including random forest, boosting, rule-based, and deep learning (neural network). The Centers for Medicare and Medicaid Services computing environment, which must be used for the 100% Medicare sample, does not provide advanced machine learning tools. Consequently, our machine learning models were based on a 5% Limited Data Set sample, which we were able to host locally. The models were developed using R software (4.2.0; available at <https://cran.r-project.org/src/base/ release 2022-04-22>). The methods we explored were the following:

1. Ranger Random Forest (RangerRF, using R package RANGER 0.13.1; available at <https://cran.r-project.org/src/contrib/Archive/ranger/released 2021-07-14>)<sup>31,32</sup>
2. Extreme Gradient Boosting (XGBoost, using R package XGBOOST 0.90.02; available at <https://cran.r-project.org/src/contrib/Archive/xgboost/released 2019-08-01>)<sup>33</sup>
3. Combination Gradient Boosting with Random Forests (XGBoostRF, using R package XGBOOST 0.90.02)<sup>33</sup>
4. RuleFit (RuleFit, using R H2O package 3.36.1.2; available at <https://cran.r-project.org/src/contrib/Archive/h2o/ released 2022-05-28>)<sup>34</sup>
5. Automated Machine Learning (autoML, using R H2O package 3.36.1.2)<sup>35</sup>

A reference model on the local system was created using logistic regression (using R package GLMNET 3.0;



available at <https://cran.r-project/src/contrib/Archive/glmnet/> released 2019-11-09).

Logistic and machine learning models were developed and evaluated on the 5% sample database using methods similar to those employed for the primary analyses, with the following exceptions:

1. The development set was restricted to 2018 admissions.
2. The set of features available for modeling were restricted to a subset of the International Classification of Diseases, Tenth Revision, variables corresponding to the 400 having the highest predictive power per endpoint as identified using Extreme Gradient Boosting (maximum depth, 4; maximum rounds, 200).
3. Separate models were developed for each endpoint, one each for unplanned and planned admissions. Sex was included among the top 400 features.

Specifics on model development for each method are detailed in the following sections.

**RangerRF.** RangerRF models (with minimum node size of 8 and maximum tree depth of 50) were identified as those with the highest out-of-bag AUC when optimized over a grid where the number of trees ranged over 1,000, 2,500, and 4,000, and the splitting method was either Gini or Hellinger.

**XGBoost.** Utilizing a 75%/25% learn/test random split of the development database, XGBoost models were identified as those with the highest test set AUC when optimizing hyperparameters over a grid where the learning rate ranged over 0.1, 0.25, and 0.4, the maximum tree depth ranged over 2, 4, and 6, the fraction of variables sampled in training each tree ranged over 0.33, 0.5, and 0.67, and the default values were used for other parameters. A maximum of 400 boosting rounds was used for each parameter combination.

**XGBoostRF.** Combination Gradient Boosting with Random Forest models were developed by boosting candidate random forest models. XGboost RF models were identified as those with the highest test set AUC when optimizing hyperparameters over the same range of values described in the prior section for XGboost, and additionally expanding the dimensionality of the grid to include 10, 20, and 40 trees.

**RuleFit.** The RuleFit algorithm creates a model in four steps. First, it fits a tree ensemble to the data. Second, it builds a rule ensemble by traversing each tree. Third, it evaluates the rules on the data to generate additional sets of features that represent interaction terms identified by the rules. Fourth, it fits a sparse Least Absolute Shrinkage and Selection Operator regression model to the enlarged pool of features containing the original 400, augmented with the newly created rule-based features. RuleFit models were identified as the resultant set of Least Absolute Shrinkage and Selection Operator models when using 400 rule-generation trees and restricting the candidate

pool of features for regression to 1,000 (*i.e.*, the original 400 plus 600 rule-based features).

**AutoML.** The automated machine learning framework H2O trains a collection of models using various boosting ensembles, random forest methods, generalized linear models, and deep learning (neural networks), with grid search to find optimal parameter values. It then uses a generalized linear model to combine the individual models into an optimal metalearner. The final models were identified as those with the highest AUC on the cross-validation test set optimized over its default parameter settings.

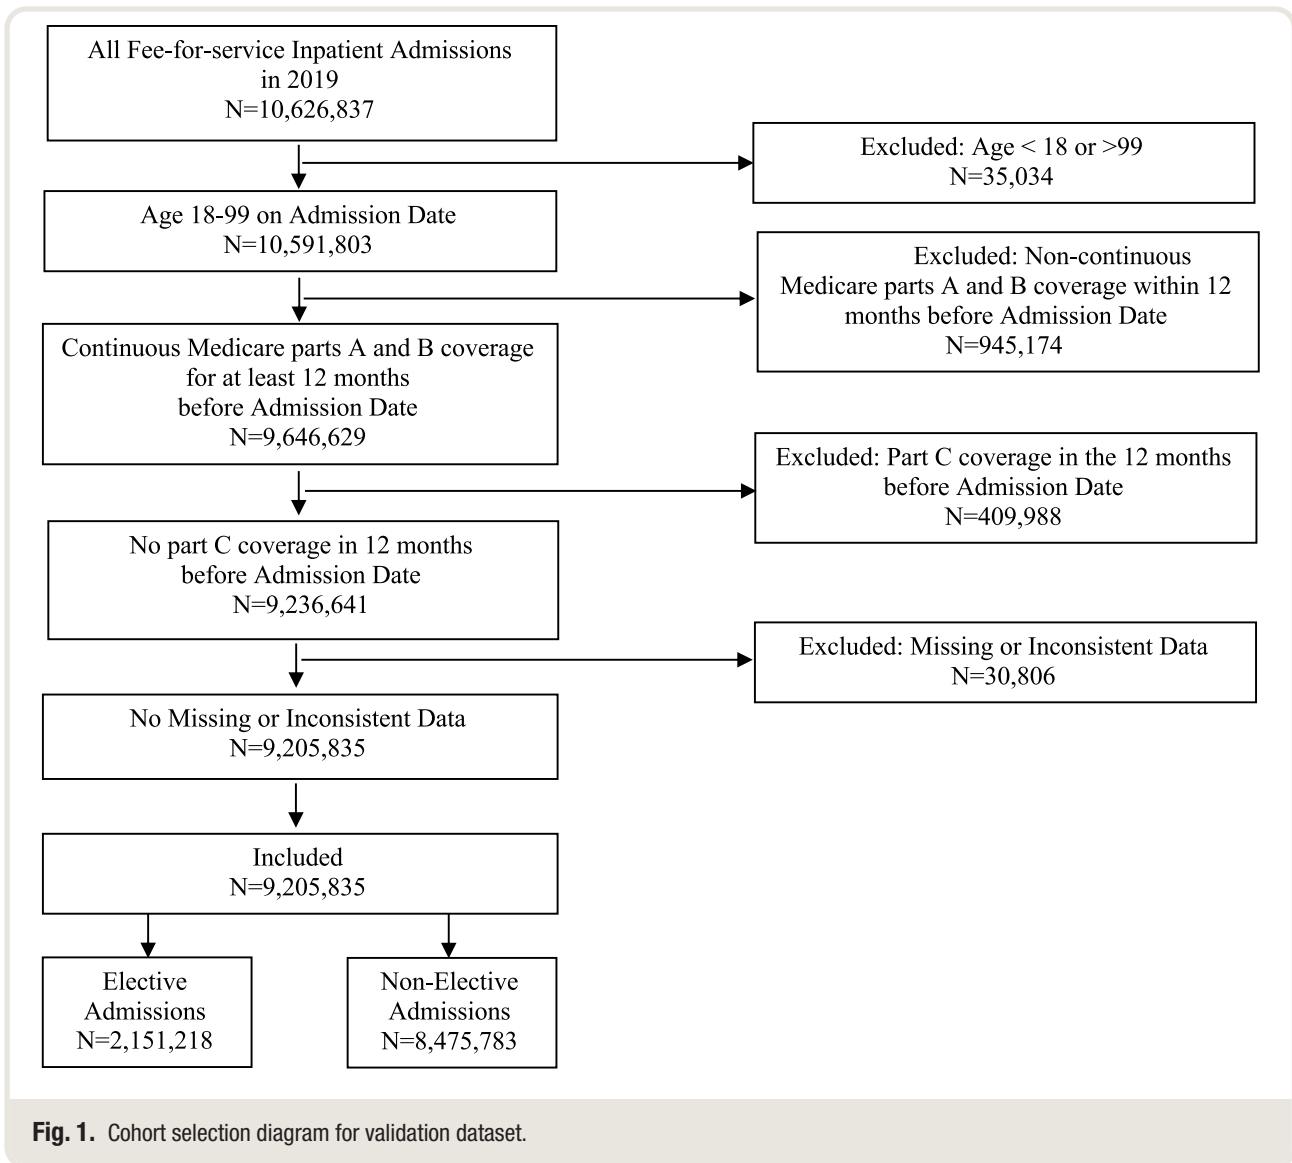
We computed performance metrics for these alternate models just as we did for our primary multivariable regression model. Due to the exploratory nature of this robustness testing component of our work, we avoided statistical comparisons among various models.

## Results

There was a total of 18,899,224 admissions in 2017 to 2018 across 9,085,968 beneficiaries in the Medicare research identifiable database who were eligible for analysis in the development set (many patients were admitted multiple times; Supplemental Figure 2, <http://links.lww.com/ALN/C923>). There were 9,205,835 admissions eligible from 5,336,265 beneficiaries for analysis in the 2019 prospective validation set (fig. 1). For our machine learning analysis, there were 476,593 admissions across 279,016 beneficiaries in the 5% limited database in 2018 for development, and 465,064 admissions from 272,220 beneficiaries in 2019 for the prospective out-of-sample evaluation. Population characteristics of the development and validation sets were similar, with a slight predominance of women (54%), and the average age was 74 yr. About 80% of admissions were unplanned (Supplemental Table 2, <http://links.lww.com/ALN/C923>).

Performance of models developed on the learn and test sets was nearly identical (not shown), confirming robustness of the modeling process and a lack of overfitting. Prospective performance characteristics of predictors of the binary events in the 2019 out-of-sample validation set are summarized in table 1. The incidence of endpoints ranged from 2.8% for in-hospital mortality to 35.9% for discharge to a care facility. The mean and range of AUCs across nine outcomes were 0.76 (0.70 to 0.82). The mean and range of the calibration goodness-of-fit measure  $R^2$  for the 99% of patients at lowest risk were 1.00 (0.99 to 1.00).

The mean and range of the observed-to-expected ratio were 0.97 (0.90 to 1.00). For the highest 5% risk population, mean and range of sensitivity were 19.2% (11.6 to 30.1%), for positive predictive value they were 37.2% (14.6 to 87.7%), and for lift (enrichment ratio) they were 3.8 (2.3 to 6.1). At the observed incidence, the mean and range of sensitivity were 31.2% (15.5 to 63.9%), for positive predictive value they were 31.2% (15.5 to 63.9%), and for lift they



**Fig. 1.** Cohort selection diagram for validation dataset.

were 3.7 (1.7 to 7.4). Note that the sensitivity and positive predictive value are equal when the detector operates at a threshold resulting in positive decisions (alerts) at a rate equaling the event incidence. A sample composite plot of charts describing performance characteristics for discharge to facility is presented in fig. 2. Similar plots for all other endpoints are in Supplemental Figures 3 to 11 (<http://links.lww.com/ALN/C923>). Excess length of stay had a root mean square error of 0.19 and  $R^2$  of 0.99.

Table 2 shows the prospective performance on the 5% limited dataset of logistic and various machine learning models developed on the 5% limited dataset and the performance of the logistic model developed on the 100% Research Identifiable File. The 5% sample appeared sufficient for comparative performance to the 100% sample because it lacked only one feature found in the 100% sample (F328, other depressive episodes.) Overall, AUC performance was similar

for each endpoint across all model types, with only minor differences in relative performance among endpoints.

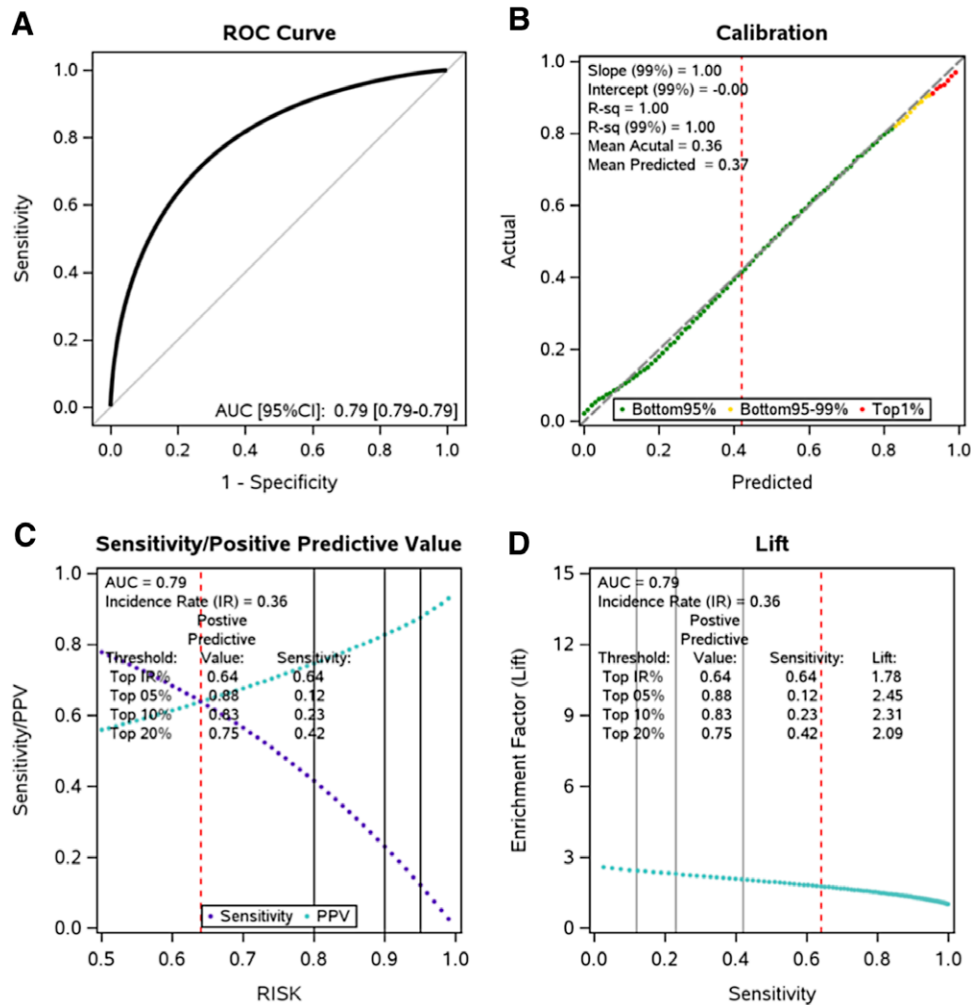
The logistic model developed on the 100% Research Identifiable File performed best on eight of the nine endpoints, probably because a larger selection of statistically significant features afforded by the larger pool of events available in the much larger database. Likewise, machine learning models developed using the larger database would presumably perform better than those developed on the 5% sample, but there is no reason to expect that relative ranking would differ much. Of models developed on the smaller 5% limited dataset, it appears that gradient boosting performed marginally better than logistic regression, but not by a clinically meaningful amount. This observation is consistent with previous work demonstrating that logistic regression provides results comparable to machine learning methods when large datasets are used.<sup>36,37</sup>

**Table 1.** Primary Model Performance on Out-of-sample 2019 Medicare Admissions

Period	Endpoint	Incidence (%)	Calibration Metrics				Discrimination Performance Metrics						
			Observed-to-Expected Ratio	R <sup>2</sup>	Correlation Coefficient*	Calibration Estimates (Intercept/Slope)	Operating Point that Alerts for Highest 5% Risk Group (%)		Operating Point that Alerts for Highest 5% Risk Group (%)		Lift	Positive Predictive Value	Lift
							Sensitivity	Positive Predictive Value	Sensitivity	Positive Predictive Value			
In-hospital	Mortality	2.8	0.95	1.00	1.00	0.00/0.94	0.82 (0.82–0.82)	30.1	17.0	6.1	20.6	20.6	7.4
90-days after admission	Discharge to facility	35.9	0.98	1.00	1.00	0.00/1.00	0.79 (0.79–0.79)	12.2	87.7	2.4	63.9	63.9	1.8
	Pneumonia	3.9	0.90	1.00	1.00	0.00/0.89	0.72 (0.72–0.72)	18.5	14.6	3.7	15.5	15.5	4.0
	Acute kidney injury	5.5	0.99	1.00	1.00	0.01/0.87	0.73 (0.73–0.73)	16.3	17.8	3.2	17.8	17.8	3.2
	Sepsis	5.8	1.00	1.00	1.00	0.01/0.92	0.73 (0.73–0.74)	18.7	21.7	3.7	20.7	20.7	3.6
	Major cardiovascular complication	6.0	0.98	1.00	1.00	0.01/0.93	0.79 (0.79–0.79)	24.1	28.9	4.8	27.5	27.5	4.6
Overall mean (95% CI)	Respiratory failure	6.3	0.98	1.00	1.00	0.01/0.90	0.73 (0.73–0.73)	20.0	25.3	4.0	23.9	23.9	3.8
	Mortality	13.4	0.96	1.00	1.00	0.01/0.95	0.82 (0.82–0.82)	21.1	56.7	4.2	44.6	44.6	3.3
	Unplanned admission	28.0	0.99	1.00	1.00	0.00/0.98	0.70 (0.70–0.70)	11.6	65.0	2.3	46.7	46.7	1.7
		12.0 (3.7–20.2)	0.97 (0.95–0.99)	1.00 (1.00–1.00)	1.00 (1.00–1.00)	0.01 (0.00–0.01)	0.76 (0.73–0.79)	19.2 (15.2–23.2)	37.2 (19.1–55.3)	3.8 (3.0–4.6)	31.2 (19.7–42.7)	31.2 (19.7–42.7)	3.7 (2.5–4.9)

Metrics include endpoint, observation window (either in-hospital or within 90-days post-admission), incidence, observed-to-expected ratio, *F*<sup>2</sup> goodness-of-fit measure, correlation coefficient, area under the receiver operating characteristics curve (AUC), sensitivity, positive predictive value, and lift (enrichment ratio). Endpoints in the table are ordered by increasing incidence within their corresponding event period.  
 \**F*<sup>2</sup> and correlation coefficient between observed and expected incidences among subpopulations using bins of 0.01 resolution of predicted risk, excluding the top 1% of subjects with highest risk. N.B. sensitivity and positive predictive value are equal when the detector operates at a threshold resulting in positive decisions (alerts) at a rate equal to the event incidence.

### Discharge to Facility



**Fig. 2.** Model performance for discharge to facility. *A*, Receiver operating characteristic curve. The curve displays the tradeoff between sensitivity and specificity over the range of possible detection thresholds. Tabulated metrics: mean and 95% CI of the area under the receiver operating characteristics curve (AUC). *B*, Calibration curve. The calibration plot displays mean actual incidence *versus* mean predicted risk of discharge to facility for populations clustered in 1% increments of the predicted risk. *Dark green*, *light green*, and *red dots* are populations of the lowest 95%, 95 to 99%, and top 1% risk. The *diagonal line* identifies the domain of ideal performance where actual and expected incidence are equal. The performance of this index is close to ideal for approximately 99% of the population. Tabulated metrics: The incidence of discharge to facility was 36%. The AUC was 0.79. Slope (99%) and Intercept (99%) are the estimates of slope and intercept of the best fit line for all subjects except the riskiest 1%. Rsq and Rsq (99%) are goodness-of-fit measures of individual results to the best fit line for all subjects and all subjects except the riskiest 1% (*i.e.*, *green dots*). *C*, Sensitivity/positive predictive value plot. Positive predictive value (*blue dots*) and sensitivity (*purple dots*) *versus* the fraction of population, sorted by the risk of discharge to facility. The *vertical red line* indicates where the number of patients above the risk threshold equals the incidence of the discharge to facility event in the population. Tabulated metrics: AUC and the incidence of discharge to facility (incidence rate). *Vertical bars* help identify the positive predictive value and sensitivity performance for detectors operating to identify the riskiest 5%, 10%, and 20% of patients. The positive predictive value and sensitivity are tabulated for these detector operating points. *D*, Enrichment factor (lift) plot. Lift (*i.e.*, positive predictive value/incidence) *versus* sensitivity. *Vertical bars* help identify the lift and sensitivity performance for detectors operating to identify the riskiest 5%, 10%, and 20% of patients. Positive predictive value, sensitivity, and lift are tabulated for these detector operating points. The AUC and incidence of mortality (incidence rate) are also tabulated.

Complication-specific models consistently identified more patients who experienced specific complications than patients selected only for mortality risk (fig. 3). Model

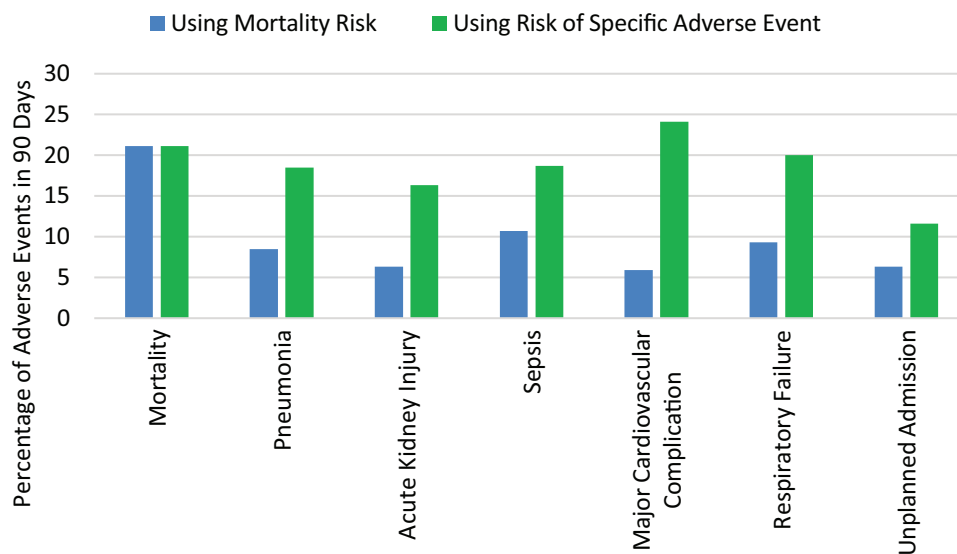
information, including open source coefficient files, is made available at <https://my.clevelandclinic.org/departments/anesthesiology/depts/outcomes-research/risk-stratification>



**Table 2.** Comparative Classification Performance across Modeling Techniques

Modeling Technique	Development Database	Endpoint and Period									
		In-Hospital					90-Days after Admission				
		Mortality	Discharge to Facility	Mortality	Acute Kidney Injury	Sepsis	Respiratory Failure	Pneumonia	Major Cardiovascular Composite	Unplanned Admission	
Logistic	5% LDS	0.808 (0.804–0.811)	0.774 (0.773–0.775)	0.817 (0.815–0.819)	0.705 (0.701–0.708)	0.713 (0.710–0.716)	0.734 (0.731–0.737)	0.690 (0.686–0.694)	0.768 (0.765–0.771)	0.700 (0.698–0.701)	
xgboost		0.815 (0.812–0.819)	0.780 (0.778–0.781)	0.821 (0.820–0.823)	0.715 (0.712–0.718)	0.720 (0.717–0.723)	0.741 (0.738–0.744)	0.701 (0.697–0.704)	0.775 (0.773–0.778)	0.702 (0.701–0.704)	
xgboostRF		0.820 (0.817–0.823)	0.786 (0.784–0.787)	0.823 (0.822–0.825)	0.720 (0.717–0.723)	0.723 (0.720–0.726)	0.745 (0.742–0.748)	0.697 (0.693–0.700)	0.778 (0.775–0.780)	0.704 (0.703–0.706)	
rangerRF		0.806 (0.803–0.810)	0.772 (0.771–0.773)	0.815 (0.814–0.817)	0.722 (0.719–0.725)	0.723 (0.720–0.726)	0.743 (0.740–0.746)	0.710 (0.707–0.714)	0.777 (0.775–0.780)	0.702 (0.701–0.704)	
autoML		0.817 (0.813–0.820)	0.782 (0.780–0.783)	0.812 (0.810–0.814)	0.710 (0.707–0.713)	0.701 (0.698–0.704)	0.725 (0.722–0.728)	0.689 (0.685–0.693)	0.767 (0.764–0.770)	0.701 (0.700–0.703)	
RuleFit		0.811 (0.808–0.815)	0.776 (0.774–0.777)	0.819 (0.818–0.821)	0.711 (0.708–0.714)	0.717 (0.714–0.720)	0.738 (0.735–0.741)	0.695 (0.691–0.698)	0.773 (0.770–0.776)	0.701 (0.699–0.702)	
Logistic	100% Research identifiable file	0.820 (0.816–0.823)	0.774 (0.773–0.776)	0.823 (0.822–0.825)	0.724 (0.721–0.727)	0.730 (0.727–0.733)	0.746 (0.743–0.749)	0.715 (0.711–0.718)	0.781 (0.779–0.784)	0.705 (0.703–0.706)	

Overall classification performance was assessed using the area under the receiving operating characteristics curve (AUC), with higher AUC values indicating better detection performance. Table cells report the AUC and 95% CI. The techniques evaluated were logistic regression (Logistic, the reference method), Extreme Gradient Boosting (XGBoost), Combination Gradient Boosting with Random Forests (XGBoostRF), Ranger Random Forest (RangerRF), Automated Machine Learning (autoML), and RuleFit (RuleFit). Two logistic models were developed (one on the 5% Limited Data Set (LDS) set, the other on the 100% Research Identifiable File). All models were prospectively evaluated on the Medicare 2019 Admissions 5% Limited Data Set.



**Fig. 3.** Percentage of adverse events in the highest 5% risk group detected using complication-specific prediction models *versus* the highest 5% risk group for 90-day mortality. *Blue bars* show the incidence of various complications in patients selected for having the highest 5% mortality risk. *Green bars* show the incidence of complications for the highest 5% risk based on complication-specific models. Use of predictors specific to adverse events results in detecting far more adverse events than using the risk of mortality alone.

(accessed October 12, 2022). The repository includes descriptions of the models and instructions how to derive the predictors from International Classification of Diseases, Tenth Revision, codes, and how to use the provided equations to make predictions.

A reasonable question is whether available medical history is a sufficient substitute for present-on-admission codes. We therefore compared the two models using codes found the year before admission *versus* performance using only codes that were *post hoc* coded as present on admission. The AUC performance for the two methods is shown for each model in table 3. Models built with historical codes were preferable, with a *P* value of 0.006. We therefore conclude that using available codes from the year before admission is superior to using present-on-admission codes—which are actually not generally available at admission since they are usually coded *post hoc*.

## Discussion

Unlike previous versions of the Risk Stratification Index, version 3.0 models are based on billing codes from the year before admission. The only information used from index admissions was admitting diagnosis, principal procedure (for planned admissions), and patient age. Consequently, our models can be implemented at admission with results immediately available to inform clinical decision-making (field tests of this approach are in progress at several institutions). Other major advances include use of International Classification of Diseases, Tenth Revision, codes, inclusion

of medical as well as surgical admissions, and many new outcomes. Risk Stratification Index 3.0 is thus a substantial advance from previous versions.

Despite restricting inputs to our models to health events captured in billing codes during the year preceding admission and limited information about the pending admission, our models performed reasonably well. AUCs in the validation set exceeded 0.70 for all endpoints, indicating satisfactory discrimination power over the range of operating thresholds. The calibration goodness-of-fit measures,  $R^2$ , exceeded 0.95 for all models, indicating strong correlation between observed and predicted values along the full continuum of risk. Furthermore, the prospective observed-to-expected ratios were between 0.95 and 1.00 across all outcomes, indicating that our models predicted outcomes well in an out-of-sample population.

We chose a conservative approach to evaluate sensitivity, positive predictive value, and lift performance by prespecifying two standard thresholds for comparing models. The top 5% risk threshold represents a means to identify subjects at the highest level of risk for a particular outcome as in a recent Artificial Intelligence Health Outcomes Challenge competition.<sup>38</sup> We also considered a threshold tethered to the incidence for each endpoint.

Observed sensitivity, positive predictive value, and lift results for each model highlight both the limitations and potential clinical application of these types of predictive models. Focusing solely on subjects with the highest 5% of risk provides only modest sensitivity of between 12 to 30% of

**Table 3.** Comparative Performance between Models Based on Historical Codes *versus* Present-on-Admission Codes Based Out-of-Sample 2019 Medicare Admissions

Period	Endpoint	1-yr Hx (No POA) [Expected Use “Historical Method”]	POA Only (No 1-yr Hx) [“POA Method”]	Pairwise Difference (Historical – POA)
		AUC (95% CI)	AUC (95% CI)	AUC Difference
In-Hospital	Mortality	0.82 (0.82–0.82)	0.84 (0.84–0.84)	–0.02
	Discharge to facility	0.79 (0.79–0.79)	0.76 (0.76–0.76)	0.03
90-days after admission	Pneumonia	0.72 (0.72–0.72)	0.67 (0.67–0.67)	0.05
	Acute kidney injury	0.73 (0.73–0.73)	0.67 (0.67–0.67)	0.06
	Sepsis	0.73 (0.73–0.73)	0.69 (0.69–0.70)	0.04
	Major cardiovascular complication	0.79 (0.79–0.79)	0.75 (0.75–0.76)	0.04
	Respiratory failure	0.73 (0.73–0.73)	0.69 (0.69–0.69)	0.04
	Mortality	0.82 (0.82–0.82)	0.82 (0.82–0.82)	0.00
	Unplanned admission	0.70 (0.70–0.70)	0.65 (0.65–0.65)	0.05
Overall mean (95% CI)	0.76 (0.73–0.79)	0.73 (0.68–0.77)	0.03 (0.01–0.05)	

Complications and resource utilization outcomes are ordered by increasing frequency within their event periods. The 1-yr historical method used available billing codes in the year before admission, and the present-on-admission method used only codes that were *post hoc* coded as being present on admission. Both methods used admitting diagnosis (not principal diagnosis) and principal procedure when appropriate. Variance is presented as 95% CI.

those who will experience each event. Although most people who had an adverse event were at lower risk levels, lift values for these top 5% riskiest patients exceeded 2.3, indicating that they were more than twice as likely as others to experience a future outcome event. The enrichment factor was particularly high for low-incidence (less than 5%) events, ranging from 3.7 to 6.1 for various models. In practical terms, this means that patients identified as being in this highest-risk category on admission are about five times more likely to experience adverse events than the general population. This 3.0 generation of Risk Stratification Index models, based solely on previous claims history and admitting diagnosis, therefore quantifies and ranks patient risk surprisingly well.

Both health trajectory and real time data will increasingly be available because Medicare is encouraging interoperability and improved access to individual claims information through Blue Button individual access,<sup>39</sup> the Beneficiary Claims Data application program interface,<sup>40</sup> and Data at the Point of Care<sup>41</sup> initiatives. Real-time availability of automatically generated Risk Stratification Index profiles and associated alerts may help reduce cognitive load for the clinician by identifying key areas of concern

in individual patients, thus potentially guiding monitoring and management.<sup>42</sup>

Although identifying patients at high risk of mortality helps to identify patients at high risk for specific adverse events associated with mortality, the use of predictors specific to adverse events results in detecting far more adverse events than using the risk of mortality alone. Accurate selection of patients at risk for specific complications therefore requires complication-specific models. A suite of predictors thus provides more information to guide risk-reducing treatment pathways in personalized care plans than overall risk of mortality.

Risk stratification models are conventionally developed from logistic regression models. Regression has the advantage of generating models that are portable and easy to deploy, apply, and interpret. Furthermore, regression models provide coefficients that identify factors that contribute most to specific adverse events, which is clinically valuable information, especially when contributing factors are modifiable. An additional consideration is that models can easily be re-run to accommodate new information obtained at a prehospitalization assessment, or even during hospitalization.

Although interpretable models are frequently preferred even among experts, machine learning models are increasingly popular.<sup>43</sup> Some reports suggest better performance for ML models compared to traditional clinical scoring systems,<sup>17,18</sup> but few have been validated on multiple external datasets. We evaluated five of the most commonly used methods. Interestingly, prediction model characteristics were similar with all six approaches—indicating that none was obviously superior and that any of the approaches is valid.<sup>11–16</sup> Therefore, an important corollary is that we appear to have identified both the value and ceiling for the amount of predictive information that can be derived from the medical claims history.

Our validation analysis was based on more than 9 million adult hospital admissions in 2019 among patients enrolled in the United States fee-for-service and Medicare/Medicaid program. Patients included in our validation represent approximately 70% of all hospital admissions in the Medicare-eligible population in the United States.<sup>10</sup> We excluded less than 0.4% of the available admissions because of missing and inconsistent values. Furthermore, data were missing nonsystematically, meaning that exclusion of these admissions was unlikely to introduce meaningful bias. Our results are therefore broadly applicable to Medicare-eligible adults. Although our 2019 sample included 1,025,099 dual eligible subjects younger than 65 yr (representing 16.3% of the 2019 dataset), our results should be cautiously generalized to younger and healthier populations.

Use of the Medicare claims database to represent individual medical histories remains controversial. For example, reliability of the Centers for Medicare and Medicaid Services registry depends completely on accurate coding. However, well-enforced federal laws promote accurate billing, and regional differences in billing appear to result from true local practice patterns rather than miscoding.<sup>15</sup> Potential errors and delays in Medicare claims processing are offset by large sample size, population diversity, and the highly structured and longitudinal nature of the dataset.

The Centers for Medicare and Medicaid Services usually uses a 30-day observation period. We used a 90-day period because previous reports suggest that 90-day outcomes may be more reliable than 30-day outcomes for measuring hospital performance.<sup>22–24</sup> Model performance at 30 days may of course differ.<sup>40</sup> A potential limitation of real-time use of the models may be incomplete or delayed access to codes in a patient's history, which may lead to underprediction of risk. However, analysis shows that predictions based on available 1-year coding history are slightly better than those based on codes for present-on-admission conditions.

In summary, we developed a suite of risk stratification models using methods similar to those used for earlier Risk Stratification Index versions. An important

distinction from previous models is that version 3.0 is based on historical information coupled with admitting diagnosis, principal procedure (for planned admissions), and patient age. Consequently, our models can be implemented at admission with results immediately available to guide clinical decision-making. Other major advances include use of International Classification of Diseases, Tenth Revision, codes rather than International Classification of Diseases, Ninth Revision, codes; inclusion of medical as well as surgical admissions; and many new outcomes. Our models predicted outcomes well in an out-of-sample population and provide clinically meaningful guidance to clinicians.

### Acknowledgments

The authors thank John Parks (Data Analyst, Health Data Analytics Institute, Dedham, Massachusetts) for his assistance in developing datasets and models. Dr. Greenwald had full access to all the data in the study and takes responsibility for the integrity of the data and the accuracy of the data analysis. He conducted the data analysis and takes responsibility for it.

### Data Sharing

The investigators' access to Medicare data is based on contracts with the Centers for Medicare and Medicaid Services (Baltimore, Maryland), which precludes sharing data. However, the data are readily available to other investigators *via* contract with the Center for Medicare and Medicaid Services. Requests for access to data to replicate these findings require a research protocol and data use agreement. For more information, contact the Research Data Assistance Center (Minneapolis, Minnesota; <http://www.resdac.org>). Statistical code will not be shared, but instructions for creating the model variables and the equations to apply provided model coefficients will be posted as described in the penultimate paragraph of the Results section.

### Research Support

Funded by the Health Data Analytics Institute (Dedham, Massachusetts).

### Competing Interests

G. F. Chamoun, N. G. Chamoun, D. Clain, Z. Hong, Dr. Greenwald, Dr. Manberg, and Dr. Jordan are employees of the Health Data Analytics Institute (Dedham, Massachusetts). Dr. Sessler is a consultant and shareholder in the Health Data Analytics Institute and is a consultant for Pacira biosciences (Parsippany, New Jersey). He serves on advisory boards and has equity interests in Calorint (Philadelphia, Pennsylvania), TransQtronics (Philadelphia, Pennsylvania), Medasense (Tel Aviv, Israel), Serenno

(Yokneam, Israel), Sensifree (Cupertino, California), Perceptive Medical (Newport Beach, California), and Neuroindex (Tel Aviv, Israel). He serves on the Board of the Foundation for Anesthesia Education and Research, and is a Senior International Fellow at the Population Health Research Institute, McMaster University (Ontario, Canada). The Department of Outcomes Research, which Dr. Sessler chairs, has research grants from dozens of companies. The other authors declare no competing interests.

## Correspondence

Address correspondence to Dr. Sessler: Department of Outcomes Research, Anesthesiology Institute, Cleveland Clinic, 9500 Euclid Ave—L1-407, Cleveland, Ohio 44195. ds@or.org. This article may be accessed for personal use at no charge through the Journal Web site, www.anesthesiology.org.

## Supplemental Digital Content

Supplemental Digital Content File, <http://links.lww.com/ALN/C923>

Supplemental Table 1: Endpoint definitions.

Supplemental Table 2: Full Medicare population characteristics.

Supplemental Figure 1: Selection of candidate diagnostic codes.

Supplemental Figure 2. Development cohort selection.

Supplemental Figure 3. Performance characteristics: in-hospital mortality.

Supplemental Figure 4. Performance characteristics: discharge to facility.

Supplemental Figure 5. Performance characteristics: 90-day pneumonia.

Supplemental Figure 6. Performance characteristics: 90-day acute kidney injury.

Supplemental Figure 7. Performance characteristics: 90-day sepsis.

Supplemental Figure 8. Performance characteristics: 90-day cardiovascular complications.

Supplemental Figure 9. Performance characteristics: 90-day respiratory failure.

Supplemental Figure 10. Performance characteristics: 90-day mortality.

Supplemental Figure 11. Performance characteristics: 90-day unplanned admission.

## References

- Lane-Fall MB, Neuman MD: Outcomes measures and risk adjustment. *Int Anesthesiol Clin* 2013; 51:10–21
- Imperial MZ, Phillips PPJ, Nahid P, Savic RM: Precision-enhancing risk stratification tools for selecting optimal treatment durations in tuberculosis clinical trials. *Am J Respir Crit Care Med* 2021; 204:1086–96
- Haas LR, Takahashi PY, Shah ND, Stroebel RJ, Bernard ME, Finnie DM, Naessens JM: Risk-stratification methods for identifying patients for care coordination. *Am J Manag Care* 2013; 19:725–32
- Levin S BS, Toerper M, Debraine A, DeAngelo A, Hamrock E, Hinson J, Hoyer E, Dunganani T, Howell E: Machine-learning-based hospital discharge predictions can support multidisciplinary rounds and decrease hospital length-of-stay. *BMJ Innov* 2021; 7:414–21
- Sessler DI, Sigl JC, Manberg PJ, Kelley SD, Schubert A, Chamoun NG: Broadly applicable risk stratification system for predicting duration of hospitalization and mortality. *ANESTHESIOLOGY* 2010; 113:1026–37
- Dalton JE, Glance LG, Mascha EJ, Ehrlinger J, Chamoun N, Sessler DI: Impact of present-on-admission indicators on risk-adjusted hospital mortality measurement. *ANESTHESIOLOGY* 2013; 118:1298–306
- Sigakis MJ, Bittner EA, Wanderer JP: Validation of a risk stratification index and risk quantification index for predicting patient outcomes: In-hospital mortality, 30-day mortality, 1-year mortality, and length-of-stay. *ANESTHESIOLOGY* 2013; 119:525–40
- Wahl KM, Moretti, E., White, W., Hale B., Gan, T.J.: Validation of a risk-stratification index for predicting 1-year mortality, 2011 Annual Meeting of the American Society of Anesthesiologists. *ANESTHESIOLOGY* 2011; pp A014
- Chamoun GF, Li L, Chamoun NG, Saini V, Sessler DI: Validation and calibration of the Risk Stratification Index. *ANESTHESIOLOGY* 2017; 126:623–30
- Chamoun GF, Li L, Chamoun NG, Saini V, Sessler DI: Comparison of an updated Risk Stratification Index to hierarchical condition categories. *ANESTHESIOLOGY* 2018; 128:109–16
- Kheterpal S, Shanks A, Tremper KK: Impact of a novel multiparameter decision support system on intraoperative processes of care and postoperative outcomes. *ANESTHESIOLOGY* 2018; 128:272–82
- Lee CK, Hofer I, Gabel E, Baldi P, Cannesson M: Development and validation of a deep neural network model for prediction of postoperative in-hospital mortality. *ANESTHESIOLOGY* 2018; 129:649–62
- Sonny A, Kurz A, Skolaris LA, Boehm L, Reynolds A, Cummings KC 3rd, Makarova N, Yang D, Sessler DI: Deficit accumulation and phenotype assessments of frailty both poorly predict duration of hospitalization and serious complications after noncardiac surgery. *ANESTHESIOLOGY* 2020; 132:82–94
- Turan A, Cohen B, Adegboye J, Makarova N, Liu L, Mascha EJ, Qiu Y, Irefin S, Wakefield BJ, Ruetzler K, Sessler DI: Mild acute kidney injury after noncardiac



- surgery is associated with long-term renal dysfunction: A retrospective cohort study. *ANESTHESIOLOGY* 2020; 132:1053–61
15. Li L, Chamoun GF, Chamoun NG, Sessler D, Gopinath V, Saini V: Elucidating the association between regional variation in diagnostic frequency with risk-adjusted mortality through analysis of claims data of Medicare inpatients: A cross-sectional study. *BMJ Open* 2021; 11:e054632
  16. Greenwald SD, Chamoun NG, Manberg PJ, Gray J, Clain D, Maheshwari K, Sessler DI: Covid-19 and excess mortality in Medicare beneficiaries. *PLoS One* 2022; 17:e0262264
  17. Hill BL, Brown R, Gabel E, Rakocz N, Lee C, Cannesson M, Baldi P, Olde Loohuis L, Johnson R, Jew B, Maoz U, Mahajan A, Sankararaman S, Hofer I, Halperin E: An automated machine learning-based model predicts postoperative mortality using readily-extractable preoperative electronic health record data. *Br J Anaesth* 2019; 123:877–86
  18. Rellum SR, Schuurmans J, van der Ven WH, Eberl S, Driessen AHG, Vlaar APJ, Veelo DP: Machine learning methods for perioperative anesthetic management in cardiac surgery patients: A scoping review. *J Thorac Dis* 2021; 13:6976–93
  19. Jing B, Boscardin WJ, Deardorff WJ, Jeon SY, Lee AK, Donovan AL, Lee SJ: Comparing machine learning to regression methods for mortality prediction using Veterans Affairs electronic health record clinical data. *Med Care* 2022; 60:470–9
  20. Collins GS, Reitsma JB, Altman DG, Moons KG: Transparent Reporting of a multivariable prediction model for Individual Prognosis or Diagnosis (TRIPOD): The TRIPOD statement. *Ann Intern Med* 2015; 162:55–63
  21. CMS: International Classification of Diseases, Tenth Revision (ICD-10). Baltimore, Centers for Medicare & Medicaid Services, 2021
  22. Hirji S, McGurk S, Kiehm S, Ejiofor J, Ramirez-Del Val F, Kolkailah AA, Berry N, Sobieszczyk P, Pelletier M, Shah P, O’Gara P, Kaneko T: Utility of 90-day mortality vs 30-day mortality as a quality metric for transcatheter and surgical aortic valve replacement outcomes. *JAMA Cardiol* 2020; 5:156–65
  23. Mizushima T, Yamamoto H, Marubashi S, Kamiya K, Wakabayashi G, Miyata H, Seto Y, Doki Y, Mori M: Validity and significance of 30-day mortality rate as a quality indicator for gastrointestinal cancer surgeries. *Ann Gastroenterol Surg* 2018; 2:231–40
  24. Damhuis RA, Wijnhoven BP, Plaisier PW, Kirkels WJ, Kranse R, van Lanschot JJ: Comparison of 30-day, 90-day and in-hospital postoperative mortality for eight different cancer types. *Br J Surg* 2012; 99:1149–54
  25. AHRQ: Clinical Classification Software Refined (CCSR), Agency for Healthcare Research and Quality, 2021
  26. AHRQ: Clinical Classification Software ICD-10-PCS (CCS), Agency for Healthcare Research and Quality, 2021
  27. ResDAC: Patient Discharge Status Table, Research Data Assistance Center (ResDAC), 2021
  28. Van der Paal B: A Comparison of Different Methods for Modelling Rare Events Data, Department of Applied Mathematics, Computer Science and Statistics. Ghent, University of Ghent, 2014
  29. Schmueli G: Lift up and Act! Classifier Performance in Resource-constrained Applications. arXiv, 2019; abs/1906.03374
  30. Bellini V, Valente M, Bertorelli G, Pifferi B, Craca M, Mordonini M, Lombardo G, Bottani E, Del Rio P, Bignami E: Machine learning in perioperative medicine: A systematic review. *J Anesth Analg Crit Care* 2022; 2:2
  31. Breiman L: Random forests. *Machine Learning* 2001; 45:5–32
  32. Wright MN, Ziegler, A.: Ranger: A fast implementation of random forests for high dimensional data in C++ and R. *J Stat Softw* 2017; 77:1–17
  33. Friedman JH: Stochastic gradient boosting. *Comput Stat Data Anal* 2002; 38:367–78
  34. Friedman JH, Popescu, B.E.: Predictive learning via rule ensembles. *Ann Appl Stat* 2008; 2:916–54
  35. H2O.ai: R Interface for H2O, 2016
  36. Cowling TE, Cromwell DA, Bellot A, Sharples LD, van der Meulen J: Logistic regression and machine learning predicted patient mortality from large sets of diagnosis codes comparably. *J Clin Epidemiol* 2021; 133:43–52
  37. Christodoulou E, Ma J, Collins GS, Steyerberg EW, Verbakel JY, Van Calster B: A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *J Clin Epidemiol* 2019; 110:12–22
  38. CMS: Artificial Intelligence (AI) Health Outcomes Challenge, Centers for Medicare & Medicaid Services, 2018
  39. CMS: CMS Blue Button 2.0, 2021
  40. CMS: Beneficiary Claims Data API, 2021
  41. CMS: Data at the Point of Care, 2021
  42. Chi EA, Chi G, Tsui CT, Jiang Y, Jarr K, Kulkarni CV, Zhang M, Long J, Ng AY, Rajpurkar P, Sinha SR: Development and validation of an artificial intelligence system to optimize clinician review of patient records. *JAMA Netw Open* 2021; 4:e2117391
  43. Rudin C: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell* 2019; 1:206–15