Paper 1361-2014

Creating a SimNICU: Using SAS® Simulation Studio to Model Staffing Needs in Clinical Environments

Chris DeRienzo, MD, MPP, and David Tanaka, MD, Division of Neonatal-Perinatal Medicine, Duke University Medical Center; Emily Lada and Phillip Meanor, Advanced Analytics Division, SAS Institute Inc.

ABSTRACT

Patient safety in a neonatal intensive care unit (NICU)—as in any hospital unit—is critically dependent on appropriate staffing. This project used SAS® Simulation Studio to create a discrete-event simulation model of a specific NICU that can be used to predict the number of nurses needed per shift. The model incorporates the complexities inherent in determining staffing needs, including variations in patient acuity, referral patterns, and length of stay. To build the model, the group first estimated probability distributions for the number and type of patients admitted each day to the unit. Using both internal and published data, the team estimated distributions for various NICU-specific patient morbidities, including type and timing of each morbidity event and its temporal effect on a patient's acuity. The final simulation model samples from these input distributions and simulates the flow of individual patients through the NICU (consisting of critical-care and step-down beds) over a one-year period. The general basis of the model represents a method that can be applied to any unit in any hospital, thereby providing clinicians and administrators with a tool to rigorously and quantitatively support staffing decisions. With additional refinements, the use of such a model over time can provide significant benefits in both patient safety and operational efficiency.

INTRODUCTION

Managing a NICU is a daunting task. The medical director is usually a senior physician who has completed medical school, three years of general pediatrics training, and three additional years of specialized training in neonatology and then has practiced as a neonatologist for a decade or more before assuming a leadership role. Similarly, a NICU's nursing director has often practiced clinical neonatal nursing for more than 10 years; usually has a master's-level degree in nursing, management, or leadership; and has significant experience in guiding smaller teams or shifts as a charge nurse or clinical leader.

Despite their years of clinical experience, nothing prepares either the nursing director or the medical director of a NICU for the administrative challenges of meeting the critical-care needs of 68 patients—all of whom are babies, ranging from hours-old premature newborns weighing less than a pound to months-old infants weighing 10 times as much as their neighbors. These challenges include assigning the right levels of acuity (a measure of how sick a patient is and thus how many resources are required to provide appropriate care) and finding the right number of nurses to staff the unit's current needs while accounting for the inevitable critical admission or patient who crashes (that is, whose clinical condition suddenly deteriorates). Then, on top of this, these administrators must plan for their NICU's future, plotting the course of growth and development for the next 5, 10, or 25 years without a data-driven understanding of today's (much less tomorrow's) needs. In the world of NICU leadership, discrete-event simulation is a game changer.

Discrete-event simulation is the modeling of a system as it changes over time, where the state of the system changes only at discrete time points. In the case of a NICU model, a discrete event might be the arrival of a baby to the unit or the transfer of a baby from the NICU to another hospital. Discrete-event simulation can be used to build a model of a complex, real-life process (like that of a NICU); then, numerically generated data from that model can be used to predict, measure, and improve system performance. Using probability distributions that are defined by a unit's existing data, today's powerful statistical software packages—like SAS® Simulation Studio used in tandem with JMP®—can take today's data and create a working model to better predict tomorrow's needs.

Although it remains impossible to account for all variables that affect a hospital unit's day-to-day variations in census (the number of patients in the unit) and acuity, it is possible to develop a model that samples from probabilistic distributions based on recent or remote past experience to simulate days, months, or years of patient admissions and discharges, as well as patient acuity, nursing needs, and transfer capabilities. Furthermore, such a model could be adjusted to predict growth and answer important business questions including the following:

- What happens to the NICU if it partners with another hospital and begins accepting 100 more transfers per year?

- How many beds will the NICU need in 10 years if growth continues at the current rate?
- How will the unit's cost structure change if nursing assistants are hired to do certain tasks, so that each registered nurse can take an additional patient and the minimum nurse-to-patient ratio drops from 1:3 to 1:4?

The possibilities are practically endless, as are the potential benefits from improved patient safety (through appropriately resourcing the unit to meet acuity needs) and from better employee satisfaction (through improved predictability of nursing assignments and reduction of stress associated with understaffing). With these myriad benefits in mind, the project team developed a model by using the Duke University Hospital NICU's physical size, structure, and nursing practices and by using de-identified, retrospective patient outcomes data as the initial inputs.

PROCESS FLOW AND MODEL INPUTS

In developing a discrete-event simulation model, it is necessary to define a process flow—called the SimNICU for this project. Each bed in the SimNICU is defined as either "critical-care" or "step-down," a designation that limits the acuity of the patient who is allowed to fill it. To match the Duke NICU's current structure, the SimNICU includes 47 critical-care beds and 21 step-down beds. In the SimNICU, patient acuity is defined using the standard nurse-to-patient ratios of 1:1, 1:2, and 1:3. That is, a 1:1 patient must be the only patient a nurse cares for during that shift, a 1:3 patient could be assigned to a nurse who has two other 1:3 patients or one additional 1:2 patient, and so on.

The functionality of the SimNICU is defined as follows. Shift changes occur every 12 hours to simulate the real nursing shift changes that occur at 0700 (7 a.m.) and 1900 (7 p.m.) in the Duke NICU. Nurses are thus reassigned every 12 hours based on a recalculation of each patient's acuity, and each nurse can take only as many patients as the acuity level of his or her most acute patient. Admissions, crashes, and deaths can occur at any point during a shift. And at the end of each shift, each patient's acuity is individually recalculated, patients are transferred out if necessary, patients are moved from critical-care beds to step-down beds if appropriate, and the nurses are reassigned.

The objects, or entities, that flow through the SimNICU represent patients (babies), and each patient entity has properties, or attributes, that include gestational age (GA), days of life (DOL) upon arrival, and acuity. The patient entities in the SimNICU are referred to as SimPatients. The number of SimPatients that arrive at the SimNICU each simulated day is sampled from a probability distribution based on historical data. After the number of arrivals has been sampled for a given day, the actual arrival times of the SimPatients are distributed uniformly across a 24-hour time period.

When a SimPatient is first generated (arrives at the SimNICU), it is randomly assigned an admission type attribute with the value Inborn (66%), Outborn-In-Network (9%) or Outborn-Out-Network (25%). Inborn patients are born at the same hospital where the SimNICU is located. Outborn patients are born at an outside hospital and transferred into the SimNICU. Inborn/outborn status is an important variable to include, because patients who are admitted from an outborn facility are significantly different from those who come from a NICU's internal delivery service in both number of admissions per day and type of patient admitted. In-network versus out-network is an important attribute for future model iterations (for example, when additional NICUs are modeled and connected to the main NICU) but has little bearing on its current function.

Next, the gestational age (GA) of the SimPatient is randomly sampled from an empirical distribution based on the admission type. Gestational age (number of weeks of gestation completed prior to the patient's birth, ranging from 22 to 42 weeks) must be included, because it serves as the primary marker of a patient's length of stay (LOS) and acuity. Babies who are born close to term and are not sick often go home within days, whereas babies who are born at 24 weeks' gestation (16 weeks before their due date) have a mandatory NICU stay of many weeks even if they do not experience any significant morbidities (long-term medical problems) of prematurity. Figure 1 shows the empirical distribution of GA (inborn admissions). Given the GA, the estimated LOS (in days) for each SimPatient is calculated as follows: if $GA < 29$ weeks, then $LOS = (37 – GA)*7$ days; if $29 \leq GA \leq 33$ weeks, then $LOS = (35 – GA)*7$ days; if $GA > 33$ weeks, then $LOS = 14$ days. Finally, the initial acuity level upon arrival for each SimPatient is set as follows: if $GA < 28$ weeks, then acuity = 1:1; if GA is between 28 and 38 weeks, then acuity = 1:2; if $GA \geq 39$ weeks, then there is a 50% chance that acuity = 1:1; otherwise acuity = 1:2 for $GA \geq 39$ weeks.

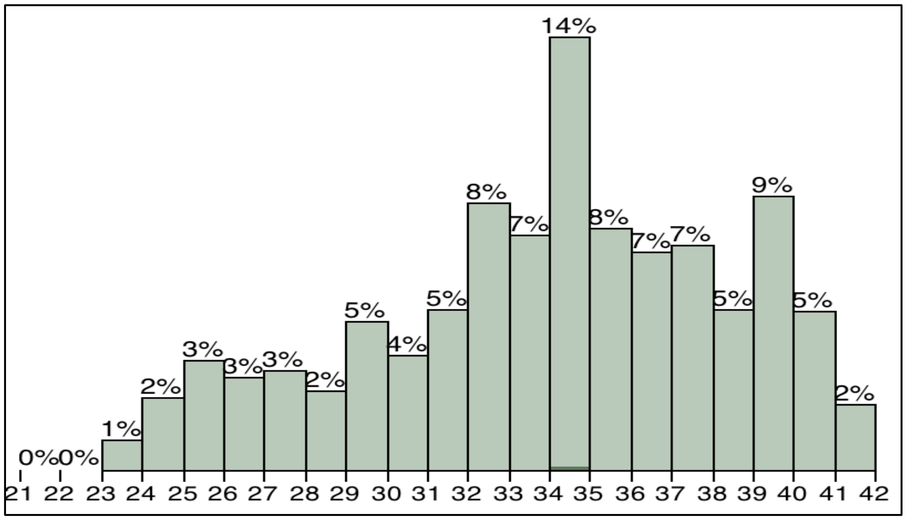

Figure 1. Gestational Age (Weeks) for Inborn Admissions

DEFINING PATIENT MORBIDITIES

Five major morbidities are known to affect the LOS of NICU patients. These clinical diagnoses include sepsis (infection in the bloodstream), necrotizing enterocolitis (NEC, a catastrophic neonatal intestinal condition), patent ductus arteriosis (PDA, a neonatal cardiovascular anomaly requiring medical or surgical intervention), retinopathy of prematurity (ROP, a condition affecting appropriate growth and development of blood vessels in the eye that leads to blindness if not promptly treated), and intraventricular hemorrhage (IVH, bleeding into the internal structures of the brain). Because the probability of each condition varies by GA, separate empirical distributions are estimated and then sampled to determine whether or not a SimPatient will experience one or more of these conditions. Furthermore, each condition occurs at different times during a patient's NICU course and affects each patient's length of stay, acuity, and potential for mortality differently. For example, NEC is several times more likely in infants born at 24 weeks than in infants born at 28 weeks, and almost unheard of in infants born at 36 weeks. Table 1 describes the basic effects of each morbidity in detail.

Morbidity	Time Window	Effect on LOS	Effect on Acuity
Sepsis	Exponentially distributed with a mean of 14.5 days	Type I (Fungi): Increase LOS by 21 days Type II: Increase LOS by 28 days Type III: Increase LOS by 14 days	Type I: 1:2 for 10 days Type II and III: No effect
NEC	Centered around DOL=30 with min=14 and max=42 days	50% mortality within 24 hours for surgical NEC; 20% mortality within 3 days for medical NEC. If no death, increase LOS by 7 days for surgical and by 21 days for medical.	For surgical NEC: 1:1 for 3 days 1:2 for 7 days For medical NEC: 1:2 for 7 days
PDA	Uniform (3,14) days	If surgical management, then increase LOS by 21 days; otherwise increase LOS by 7 days.	If surgical management: 1:1 for 3 days 1:2 for 7 days Otherwise, no effect
ROP	(32 – GA)*7 days (Occurs at 32 weeks)	Stage I: No effect Stage II: Increase LOS by 21 days Stage III & IV: Increase LOS by 42 days	Stage I & II: No effect Stage III & IV: 1:2 for 3 days when reach 32 weeks old
IVH	Uniform (7,14) days	Increase LOS by 10 days	None

Table 1. Defining Morbidities for Individual Patients

To estimate the various probability distributions and other inputs required for the SimNICU, the team accessed an existing database that includes infants who were born from January 2008 through June 2013. This database is maintained by the Duke NICU's administrative staff and is accessible to the unit's leadership team. It contains demographic information, diagnoses (including date of diagnosis for specific disease states or therapies), billing and coding details by day of stay, and discharge status (sent home, transferred to another unit, or deceased). With the approval of the Duke Institutional Review Board, the data were de-identified, analyzed, and in some instances combined with published national-level outcomes data (Stoll et al. 2010) to estimate the admission and morbidity inputs for the SimNICU. The final probability distributions that the model uses are anonymized; no SimPatients are directly based on any actual Duke patient.

DETAILED PROCESS FLOW

With the use of the anonymized admission and morbidity distributions, each SimPatient's entire course on admission to the SimNICU is predefined, including if and when a crash will occur (to the hour), shift-to-shift variations in acuity, and whether or not a SimPatient will survive to discharge or transfer. Upon arrival at the SimNICU, each SimPatient is assigned an initial LOS based on GA. If the SimPatient crashes at some point while in the SimNICU, then the LOS might be increased. After each simulated day, the LOS for each SimPatient is reduced by one day. A SimPatient exits the SimNICU in one of the following ways: by succumbing to one of the morbidities identified on admission, by dying from another aspect of prematurity that is not crash-related, or by surviving to discharge or transfer. The specific criteria for discharge and transfer are shown in Table 2. In this table, DOL (days of life) for a SimPatient is the number of days since birth, and the post-menstrual age (PMA) is computed as follows (in days): PMA = GA*7 + DOL. Discharge can occur at any shift change when a SimPatient meets the criteria. A transfer, however, requires that a SimPatient meet individual criteria, that the SimNICU meet the census-based requirement for transferring patients, and that there is a transfer bed available. The number of available transfer beds is sampled from a probability distribution at the start of each day. At the end of a shift, if a SimPatient has not crashed and its PMA is greater than or equal to 210 days, then the acuity level is set to 1:3. Figure 2 shows a high-level process flow diagram for an individual SimPatient through the SimNICU from admission to exit.

Mode of Model Exit	Necessary Conditions
Discharge	Acuity = 1:3 and remaining LOS = 0
Transfer	If GA < 29, then PMA ≥ 238, acuity = 1:3, and no ROP If 29 ≤ GA ≤ 31, then PMA ≥ 224, acuity = 1:3, and no ROP If GA ≥ 32, then DOL ≥ 7, acuity = 1:3, and no ROP

Table 2. SimPatient Criteria for Discharge or Transfer

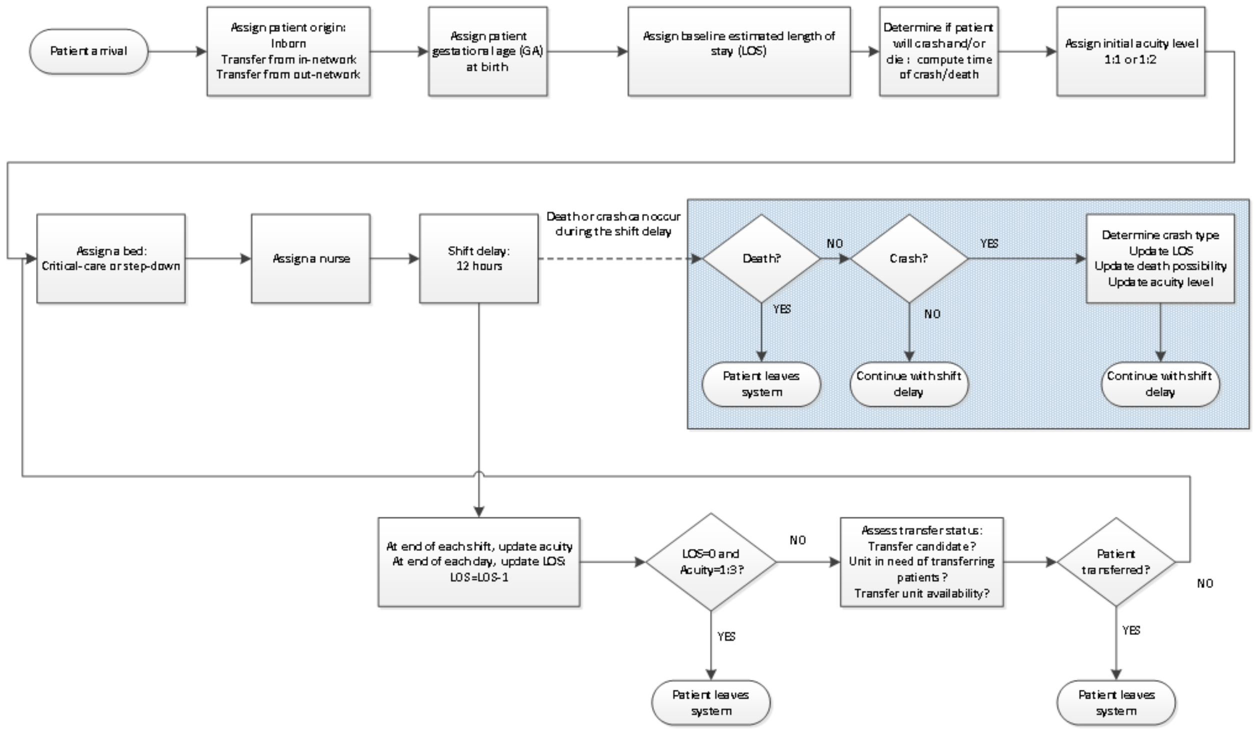

Figure 2. SimNICU Process Flow Diagram

OVERVIEW OF SAS SIMULATION STUDIO MODEL

The process flow in Figure 2 was translated into a discrete-event simulation model by using SAS Simulation Studio. Figure 3 shows a high-level view of the model in SAS Simulation Studio. Each yellow square in the model is a compound block that can be opened to display the Simulation Studio logic contained within. Figure 4 shows the expanded AssignNurses_ICN compound block. For additional information about SAS Simulation Studio, see the SAS Simulation Studio documentation (SAS Institute Inc. 2013).

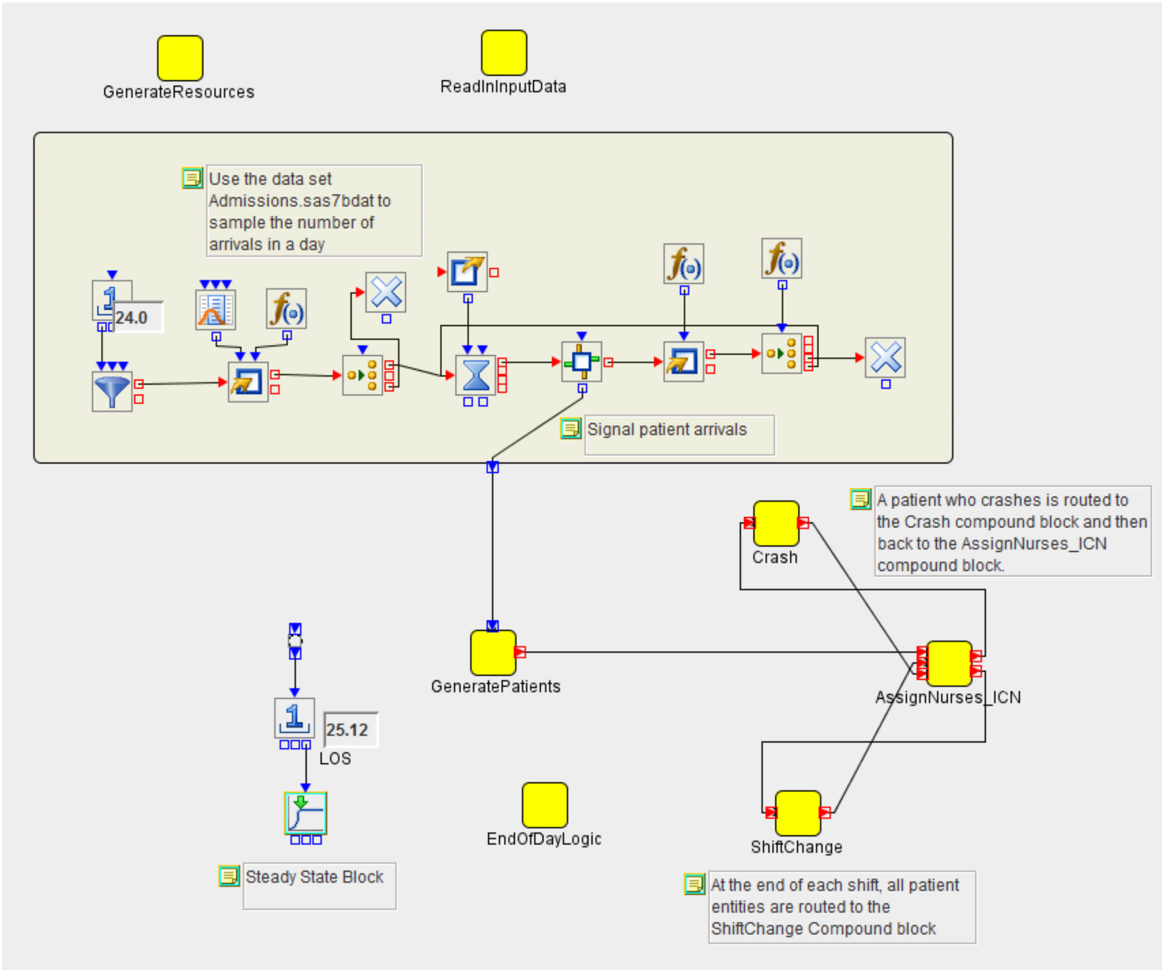

Figure 3. Screenshot of Final Model in SAS Simulation Studio

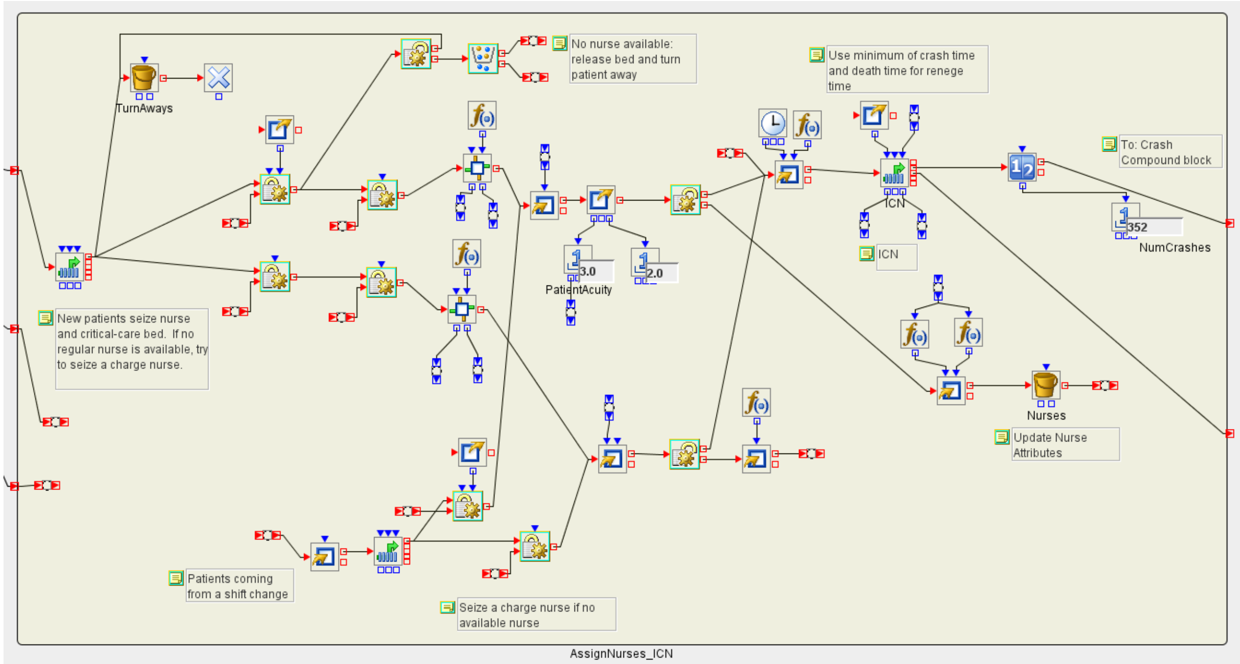

Figure 4. Expanded Simulation Studio Logic for Assigning Nurses

RUNNING THE MODEL

The following response data are collected by the SAS Simulation Studio model, using a simulation run length of one year: number of admissions, number of admissions with a GA less than 28 weeks, total number of transferred patients, and total number of patients who die. In order to study the effect of available resources on these responses, the numbers of nurses, step-down beds, and critical-care beds are used as factors in the model.

Additional responses of interest include the average daily census (the average number of patients in the SimNICU at the end of each day), LOS of all patients, and LOS of all patients who have a GA less than 28 weeks. Unlike the other responses for this model that are based on a simulation run length of one year, these are long-run performance measures, and steady-state analysis methods are used to estimate them. In particular, the SAS Simulation Studio Steady State block (bottom left of Figure 3) is used to obtain statistically valid confidence intervals for the average daily census and the average LOS. When the Steady State block is used, a simulation end time is not specified, and the model runs until enough data are collected for the response of interest to generate a confidence interval.

It is also necessary when running this model to consider system warm-up. Because it is impossible to start this simulation in steady-state operation (which would require knowing not only how many patients to prepopulate the model with but also the current state of each patient), the model starts empty and idle (with no patients in the system), and it takes some amount of simulated time before the system reaches a point that resembles steady-state operation. Whether the model is being run for one simulated year (the terminating case) or for a long period of time to estimate steady-state performance measures (the nonterminating case), it is important to accurately estimate the length of the warm-up period and include in the analysis only observations that are generated after the end of the warm-up period. If observations that are generated before the end of the warm-up period are used, then the resulting response point estimates might be biased.

For the nonterminating case, the Steady State block has a built-in algorithm for computing a confidence interval for a specified response that includes a method of detecting the end of the warm-up period and truncating the data so that the delivered point and confidence interval estimators will not be contaminated by initialization bias. The output from the Steady State block includes an estimate of the warm-up length, and this estimate can then be used as a statistic and data clearing time when the model is run for the terminating case. For responses that require a simulation run length of one year, the specified simulation end time is actually one year plus the length of the warm-up period. When the end of the warm-up period is reached during the simulation run, all collected data are cleared for the responses of interest, and only observations that are collected after the end of the warm-up period are included in the analysis. For the case where the number of nurses is 30, the Steady State block warm-up length estimate is 64 simulation days. To be conservative, a warm-up length of 120 days and a simulation run length of 485 days are used to estimate the responses for the terminating case.

MODEL VALIDATION AND RESULTS

To validate the model, the averaged results that are obtained from 50 independent replications of the terminating SimNICU model (data collected for 50 replications of the simulation, with a run length of 485 days each) are compared to the averaged actual data from the Duke NICU over a five-year period covering 2008–2012. Table 3 shows these results, including 95% confidence intervals (CIs) for the estimated mean response. Note that the SimNICU estimates of the average daily census and the length of stay are the results computed by the Steady State block in SAS Simulation Studio.

Response	Model with Nurses=30 Mean (95% CI)	Model with Nurses=26 Mean (95% CI)	Actual Mean (95% CI)
Admissions in 1 year	862.96 (855.45, 870.47)	817.58 (809.32, 825.84)	792.2 (732.4, 851.0)
Admissions for GA < 28 Weeks in 1 year	130.16 (127.92, 132.40)	122.72 (120.87, 124.57)	119 (108.9, 129.1)
Total Transfers in 1 year	281.7 (272.98, 290.42)	213.6 (205.83, 221.37)	255.2 (168.6, 341.7)
Total Deaths in 1 year	33.1 (31.43, 34.77)	30.46 (28.75, 32.17)	38.4 (33.7, 43.1)
Average Daily Census	59.61 (59.23, 59.99)	55.93 (55.67, 56.22)	57.1 (53.5, 60.7)
Length of Stay (days)	25.62 (25.26, 25.97)	27.16 (26.81, 27.58)	26.3 (25.1, 27.7)
Length of Stay (days) with GA < 28 Weeks	79.98 (78.71, 81.85)	81.65 (78.84, 83.02)	86.0 (81.1, 90.8)

Table 3. Model Predictions versus Actual Data

In theory, the Duke NICU is designed to operate with 30 nurses. However, there are often fewer than 30 nurses in the NICU during any given shift (because of variations in scheduling needs, absences due to illness or emergency, and so on). Consequently, it is difficult to compare the simulation model results from using 30 nurses with the actual results. To compensate for this discrepancy, the model was also run using 26 nurses; the results are shown in Table 3.

DISCUSSION

This project sought to address two main business intelligence needs within NICU management: (1) the need for an accurate method of modeling patient mix, patient acuity, staffing needs, and costs in the present state; and (2) the need to forecast how changes in a unit's physical structure, staffing, referral patterns, or patient mix will affect the NICU in a future state. The team's first goal has been met, as evidenced in Table 3, with a model that accurately and predictably reflects the actual data from the Duke NICU and its real patients over a recent five-year period. This critical first step provides an unprecedented snapshot of the present state for NICU managers: even before varying the inputs, this present-state model clearly identifies which levers managers can control and provides an opportunity to bolster resource requests from senior management with data-driven projections.

Although the present-state model alone has significant value for NICU managers, it is in meeting the second goal—modeling the future state—that the benefits increase exponentially. To give just one example, to build Table 3 the team modeled a unit staffed by 26, 27, 28, 29, and 30 nurses per shift; the results for 26 and 30 nurses are shown in the table. The Duke NICU's actual staffing averaged closer to 26 per shift during the period that was analyzed; thus, the 30-nurse model demonstrates how the addition of 4 nurses per shift would affect daily census (55.93 to 59.61) and yearly census (817.58 to 862.96). The value in such information, combined with a unit's proprietary cost information, becomes immediately clear as a NICU's management contemplates the return on investment for additional hiring per added patient and how to staff the unit for growth. With further variation in the inputs, a manager can now answer in a cost-effective manner how many additional nurses to hire for each bed added during a growth phase, as well as how variations in bed allocation (such as building more step-down beds versus adding critical-care beds) affect the flow of patients over a week, a month, a year, or a decade.

It is important to note that although the base model is built on the Duke NICU, the clinical data that underlie the calculations of length of stay and morbidity frequencies and the likely rough estimates of admission frequencies for various gestational ages would be applicable to most level-3 or level-4 inborn NICUs across the country. Thus, with modification of a handful of critical variables unique to an individual NICU—number of critical-care and step-down beds, number of nurses, transfer pattern, and actual admission probability envelopes—this model is easily transferrable to other units.

Furthermore, the basic discrete-event simulation structure that underlies the model is applicable to almost any clinical area in any hospital that has access to the right data. Most clinical managers already have some understanding of the major drivers of admission, discharge, staffing, and length of stay within their units, and they probably either manage these data themselves or have access to the data through their hospital's performance improvement department or other business management office. Indeed, the team dedicated the first several weeks of this project to identifying and describing these drivers as clearly and accurately as possible. This process alone uncovered the answers needed to build the clinical portion of the Duke NICU model, using both published outcomes data and existing administrative data.

With the clinical data in place, the project team then spent several months refining the process diagram and its underlying model logic around the unit's shift-to-shift and day-to-day operation. It is this logic that represents the most valuable, transferrable product of the team's efforts: with additional modification and testing, the "guts" of this work can be used to provide leaders in any clinical unit, in any hospital, anywhere in the United States with better business intelligence.

CONCLUSION

NICU patient mix and acuity can be reasonably and accurately modeled over time by using discrete-event simulation and a combination of internal and published data. With the right inputs, the proposed method can be applied to any hospital unit and represents an opportunity for unit leadership to quantitatively determine staffing needs and improve patient safety.

REFERENCES

SAS Institute Inc. 2013. *SAS Simulation Studio 13.1: User's Guide*. Cary, NC: SAS Institute Inc.

Stoll, B. J., Hansen, N. I., Bell, E. F., et al. 2010. "Neonatal Outcomes of Extremely Preterm Infants from the NICHD Neonatal Research Network." *Pediatrics* 126:443–456.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the following author:

Chris DeRienzo, MD, MPP

Before July 1, 2014:
Enterprise: Division of Neonatal-Perinatal Medicine, Duke University Medical Center
Address: DUMC Box 2739, Durham, NC 27710
Work phone: 919-681-6014
Fax: 919-681-6065
E-mail: chris.derienzo@duke.edu

After July 1, 2014:
Enterprise: Chief Patient Safety Officer, Mission Health System
Address: 509 Biltmore Avenue, Asheville, NC 28801
Work phone: 828.213.3517
Fax: 828.213.0196
E-mail: chris.derienzo@msj.org