Building a Smarter Staffing Model

NO COMMENTS

Better analytics can help your emergency department staff for tomorrow's peaks and valleys

Better analytics can help your emergency department staff for tomorrow's peaks and valleys

"Why is nurse shift-change at 6:00 AM?" I asked the Nurse Manager.

"Well there's no parking by 7, so we try to get in ahead of the rush."

"Oh. Well there are very few patients here at 6— volume doesn't really pick up until 8."

(Pause) "But there wouldn't be any parking at 8."

As a process improvement consultant, I have conversations like this all the time. Despite the enormous impact that staffing and scheduling have on a department's bottom line and patient

JOB OPPORTUNITIES

See all jobs | Post a job

RELATED ARTICLES

The Mask Debate
By SALIM R. REZAIE, MD & WILLIAM SULLIVAN, DO, JD

Projecting the Future of Emergency Physician Workforce
By GILLIAN SCHMITZ, MD, FACEP

It's Not Just About Wearing Masks
By SALIM R. REZAIE, MD

LATEST TWEETS

Tweets by @epmonthly

Emergency Physicians Monthly Retweeted

Mark Ramzy, DO, EMT-P
@MRamzyDO

@UMEmergencyMed's @CritCareguys on Crashing Pts at #ResusX21

💔Obese pts have no cardiopulmonary reserve, apnea time ≅ 4 min
💉Dose patients correctly, weight use is drug specific(👇)
💨Adjust Ventilator (higher PEEP & RR)#FOAMed @EMNews @epmonthly @smuramed @srrezaie @doc_mg

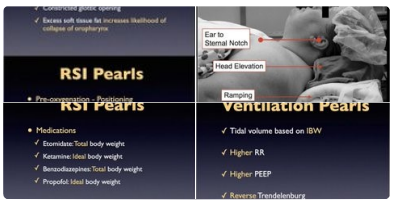

Nov 2, 2021

Embed View on Twitter

throughput, these conversations have made it clear to me over the years that most EDs have a large disconnect between tactical staffing and scheduling activities and high level strategic objectives related to budget and performance.

So why does this happen? The answer is simple: staffing and scheduling are inherently complex. It is very difficult to confidently say how decisions made about staffing and scheduling will impact strategic objectives. As a result, EDs often rely on gut feelings and corrective actions in lieu of hard data. And in a world without hard data, arguments based on flimsy premises like the availability of parking are actually quite viable.

Fortunately, the increasing availability of data and advances in how that data can be analyzed are enabling departments to evaluate staffing and scheduling decisions within the broader context of overall department strategy. The widespread adoption of Electronic Medical Records (EMRs) has dramatically increased healthcare data availability at the same time that advances in the fields of data science, optimization and simulation modeling have drastically increased the value that can be extracted from these new datasets. For the first time ever, EDs are in a position to fully understand the implications of staffing and scheduling activities, leading to incredible opportunities to capture cost savings, improve performance and boost employee morale. And when improvement opportunities are supported by objective, data-driven analysis, collective resistance to change is much easier to overcome.

Developing a "Staffing Strategy"

In the future all EDs will view staffing and scheduling as a strategic issue, and manage these activities with a Staffing Strategy. A Staffing Strategy defines an achievable set of performance and budget goals, supports them with staffing and scheduling policies, and provides clear direction for the future in the form of a concise, multiyear plan that covers a range of feasible scenarios. These strategies guide scheduling activities in a way that directly support performance goals while minimizing total manpower requirements and hours worked. An effective strategy takes into account the complexity of an ED, and is substantiated by objective, data-driven analysis. Without a staffing strategy, tactical staffing and scheduling activities are disconnected from departmental strategic goals, guided by reactive "guesstimation," and considered in relative isolation.

A staffing strategy can be developed following a process that progresses through several staffing and scheduling activities that departments already perform. However, to remove the uncertainty associated with making changes to such a complex system, departments can replace guess-and-check approaches with data science and advanced analytical techniques. With objective data-driven analysis, department leadership defines clear sets of feasible goals supported by evidence-based staffing and scheduling policies, and an actionable long-term plan.

1. Document Current State Planning

A thorough review of existing performance and budget goals is the first step in defining a staffing strategy, as a clear understanding of both is important to understanding if the two are at all compatible. For example, your facility may have a median length-of-stay goal of two hours, and a staffing budget that allows for 72 physician hours per day. Is it even possible to achieve a two hour length-of-stay with only 72 physician hours per day? Or conversely, is 72 hours more than what is actually needed to achieve a two hour length-of-stay? Simulation modeling techniques can provide answers to these kinds of questions (more on this technique later). If the goals are not compatible, developing a staffing strategy will facilitate adjustments.

2. Model and Predict Future Demand

The next step in creating a staffing strategy is predicting future demand. All departments use some

sort of demand prediction to inform schedules; the more accurate your demand predictions are, the better you can schedule. The ability to predict demand far into the future facilitates the ability to create long term plans for managing your manpower.

The simplest approach to demand prediction is to calculate the average number of patients that present each hour of the day. Many EDs used to do this, but are increasingly looking at the 80th or 90th percentile of demand based on the false premise that staffing to the average will leave them understaffed half the time. While this is not true, it is a very common misconception. First, remember that the average is not the same as the median: when it comes to arrival patterns whose distributions are highly skewed, the average and the median are generally not even close to one another. If you predict the median you can expect to under-predict half the time, but this does not tell you anything about the average arrival rate.

Regardless, the reason that EDs who staff to the average often feel understaffed is because they often apply flawed logic when translating average patient arrive rates to staffing levels. This logic commonly goes like this: if the average hourly patient arrival rate is 10 patients, and if physicians see 2.5 patients/hour, then simply divide 10 by 2.5 and determine that you need 4 physicians. If those 10 patients showed up in perfectly spaced 6 minute intervals and other resources like beds, nurses and ancillary services weren't constrained, this approach would work very well. However, ED patient arrival patterns vary wildly. This means that patients tend to present in clusters followed by periods of relative calm, and then more clusters. Every triage nurse sees this pattern occur daily. The amount of variance changes from department to department, and even within EDs from day-to-day and hour-to-hour. Knowing the variance is as important as knowing the average. In fact, it can be shown that if the variation in arrival rates doubles in a simple queue, waiting times will increase by approximately fifty percent.

Data science techniques can help EDs understand not only how many patients might arrive on average, but also how much variance there might be in those arrivals. But what is "data science" actually? For this application, these techniques amount to performing regression on patient arrival data. You could do a simple version of this in a spreadsheet. However, unlike the ordinary least squares regression available in tools like Excel, more advanced models are able to account for the possibility of high variance (sometimes called overdispersion by statisticians). The details of these models are outside of the scope of this article, and truly, a skilled data scientist is needed to select the best model, and to tune it to a particular ED's dataset. At ReefPoint Group our team will generally use quasi-Poisson, negative binomial, and zero inflated regression models which consider factors like hour-of-day, day-of-week, month, and season to predict demand. When patterns are found in the data that are not explained by these models, our team identifies additional factors such as special events, holidays or long term growth trends to account for variance and refine predictions. We will also segment patients into subpopulations to look for any patterns that occur for certain age groups, acuity levels or presenting conditions. For example, we often find that events like high school football games can drive ED demand for pediatric patients, while home NFL games can drive demand for older "children." For some EDs this demand can be significant enough to warrant an additional provider. Conversely, we often also find factors that limit demand such as specific holidays. In these cases, decreased demand may be significant enough to go from triple coverage to double.

3. Define Scheduling Targets

From demand predictions you can create scheduling targets, which are the number of staff by type you need on your schedule each hour of the day in the future. Scheduling targets can be calculated by simulating the ED process and determining how busy resources are throughout the day (called resource utilization). Arrival patterns certainly drive resource utilization (and must be

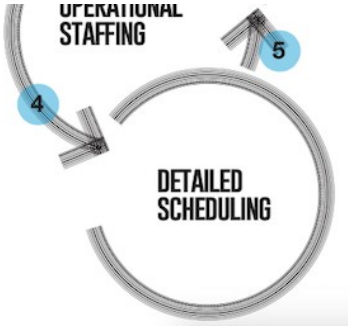

accurately modeled), but it is also impacted by many other factors such as your process and the availability of other resources (e.g., ancillary services). For example, although many patients may be arriving to your ED in a given hour, if you are in bed lock because of boarding patients, your providers will not be utilized!

Setting scheduling targets requires department leadership to define policies about the level of patient demand they will staff to. For example, your prediction may say that most likely, three providers will be needed. However, there is a 20% chance that four will be needed and a 5% chance that five will be needed. Whether you staff three, four or five providers, this decision should be informed by a clear policy that defines the way the department chooses to balance the risk of understaffing with the cost of overstaffing. Simulation modeling can provide insights into how departments will perform while staffed to various demand levels to inform this decision.

However, in the end, departments must accept that some fraction of the time staff will be playing candy crush, and some fraction of the time staff will be overwhelmed because no demand predictions or simulations are perfect. How often each of those scenarios might occur is directly relatable to staffing budget and has implications on overall performance.

4. Analyze Constraints, Define Shifts

With scheduling targets, departments can now use advanced analytical models to find the best way to allocate staff against those targets. The result of this step is to create scheduling templates which can be passed onto managers who build daily schedules. At this point in the process it is beneficial to analyze scheduling constraints, which are any rule that limit how you schedule staff. The most common scheduling constraints are shift length and shift start times. An example of a scheduling constraint are provider shifts being only 10 hours long. It is easy to see how constraints can lead to higher costs; for example, if you need a provider for only 5 hours to cover patient demand but you have to schedule them for 10, you end up with waste.

Imagine if you knew that not allowing providers to work 8 hour shifts costs your department $300,000 a year. While providers may not be happy about working 8 hour shifts, you may decide that the potential savings are significant enough to change your scheduling policies (in cases like this departments often pass a fraction of the savings onto staff by paying them a small premium for working ‘undesirable’ shifts).

Constraint analysis based on optimization techniques provides this kind of information. In fact, modern techniques are able to determine the cost associated with each and every one a department’s scheduling constraints. For example, in addition to knowing the savings gained from 8 hour shifts, these techniques will tell you the cost savings from shifts ranging in length from 1 hour to 12 hours. By reviewing all constraints together, departments are able to pick and choose among various policy configurations, and land on the one that best balances practical concerns like employee morale against dollars and cents before producing scheduling templates.

All constraints cost you money and are worthy of review. Yet these constraints are often not immediately obvious as they are not hard-and-fast rules, but rather long standing practices or conventions. Have you been wondering about those nurses I mentioned, starting at 6 AM so that they could find parking spots? This was not a policy that was written down anywhere—however, it was absolutely a constraint. In fact, after analysis of the department that inspired that story, my team determined that this constraint added fifteen percent to the total cost of nurse staffing! With evidence like this, it became very difficult for the department to not challenge this long standing practice. But don’t worry, those nurses are still getting prime spots right outside the building. While

the schedule could have been made more efficient, it quickly became clear that parking was the major morale issue for the nursing staff. As a result, the department opted keep their 6 AM start time and focus on a combination of the other opportunities to realize cost savings.

5. Finalize Staffing Strategy

Once you've documented goals, predicted demand, analyzed constraints, and defined shifts, the results of this process can be fully implemented by documenting your findings as a formal staffing strategy. The combined insights of all of these steps will help you determine what a feasible set of goals are, the best combination of staffing and scheduling policies for your department, and predict how these may change overtime. It is generally accepted as good practice that you revisit this document at least every six months to ensure you are on track with your long term plans.

David Gilinson *is a Senior Management Consultant and data scientist at ReefPoint Group, an analytic consulting firm headquartered in the Annapolis, Maryland*

Editor's Note: Death by Data

I can certainly appreciate that data science, when applied to scheduling, can give ED leadership a wide range of options to meet patient demand. Who wouldn't want a clear strategy to define what degree of overstaffing and understaffing is sought? And EPs should be comfortable working odd hours with uneven demand to get the job done, keep patients happy and healthy, and help the department meet its goals.

But when consultants are helping ED leadership develop their staffing strategy, I hope they all remember: staff is made up of people. Data science may suggest all kinds of novel work schedules that end up making parking – or child care, or advancing one's education – difficult. Emergency medicine would be wise to watch this unfold in other industries. The New York Times recently profiled a Starbucks barista whose relationships and college plans were thwarted by a convoluted, data-driven schedule.

Most EPs are fortunate to have the resources to meet the challenges of new schedules designed by modern data science. But not everyone in our department is so lucky. Hopefully department leadership has enough of a grasp on staff morale and needs — call it "old-school scheduling strategy" – to know when the algorithm needs to be chucked out the window.

Nicholas Genes, MD, PhD *is a senior editor at EPM and a clinical informaticist at Mount Sinai Medical Center*