


Predictive Analytics for Capacity Planning:

All Forecasts Are Not Created Equal



Introduction

The use of predictive analytics to manage healthcare provider capacity is one of the newest applications of this powerful tool. Faced with unprecedented pressure to produce better outcomes at a lower cost, healthcare organizations must figure out how to achieve sustainable cost reductions exceeding 20%¹ across their system, while improving quality, safety and outcomes. With labor costs typically comprising up to 60% of total expenses², one of the best ways to reduce operating costs is to align staffing and capacity to predicted patient demand.

Whether a healthcare system is experiencing variable, declining, steady or increasing volumes, rightsizing staffing to demand is critical to operational performance.

Because patient demand forecasting is so new, few organizations have sufficient exposure to full-spectrum

systems to rank various offerings in terms of the value they deliver. If you've never had anything to work with but annual budget staffing ratios, even a static forecast that predicts census over the next few weeks or months may seem like a big improvement. Granted, it's better than no forecast at all. However, if your goal is to sustainably

With labor costs typically comprising 60% of total expenses, one of the best ways to reduce operating costs is to align staffing and capacity to predicted patient demand. Whether a healthcare system is experiencing variable, declining, steady or increasing volumes, rightsizing staffing to demand is critical to operational performance.

reduce operating expenses by transforming how you allocate staff and resources, you need to understand the full spectrum of forecasting power that it will take – and discern the difference between vendor claims and reality. This guide will show you what to look for.

What Is Analytics?



Analytics is an academic and business discipline that has been around for decades. Broadly speaking, it is the use of computing to mine data to improve performance. Common examples of applied

analytics include customer segmentation, voter redistricting and financial portfolio balancing.

In healthcare, analytics has long been used to provide financial decision support and for long-range strategic planning. Today, analytics is used to aggregate data across care settings and practitioners, compare financial and clinical outcomes, determine the cost and value of care, and reduce variation in practice.

Recently the focus has shifted to using analytics to address operational issues and process improvement. The pressure to meet the Institute of Healthcare Improvement's Triple Aim³—to improve population health, improve the patient experience and reduce cost per capita—requires a broader

strategy. Although high-level strategic initiatives related to health reform can contribute to process improvement, their impact will fall short without tools and tactics that focus on optimizing day-to-day operations and resource allocation.

Real-Time Analytics

While retrospective analytics can be powerful, by definition it looks only at what happened in the past to identify patterns and trends. To provide leaders with information that facilitates proactive decisions, systems need to provide continuously updated operational information.

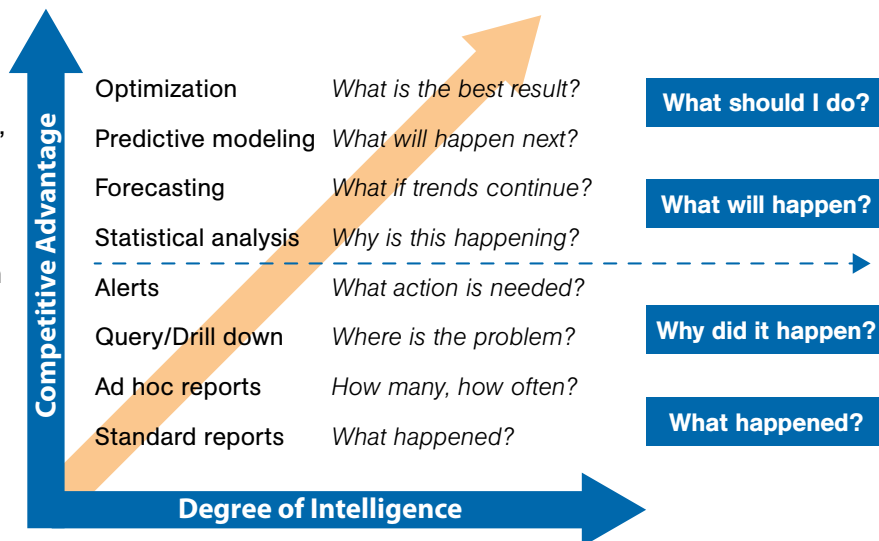
The shift to real-time analytics considerably increases the value proposition. For example, power companies use real-time analytics to monitor current energy demand across a regional grid and adjust supply to avoid brown- or blackouts. Air traffic controllers monitor prevailing winds and other conditions to dynamically adjust flight paths and help ensure on-time arrivals. Even your car dashboard provides real-time information on speed, fuel and engine status.

Predictive Analytics

Predictive analytics takes data harvesting to the next level. It uses available data not only to determine patterns, but also to forecast future outcomes with reliable precision. The most obvious example is weather forecasting, which has become markedly more accurate in recent years, despite increasing variability in weather patterns. Consumers have predictive analytics to thank for the transition from mass marketing to eerily personalized online shopping experiences. In retail and manufacturing, predictive analytics helps businesses lower costs by monitoring exactly when they will run out of each enabling “just in time” reordering. Analytics guru Tom Davenport depicts the road to analytics maturity in Figure 1.

Figure 1

Paradigm Shift Required for Analytics Maturity



Source: *Competing on Analytics*, Tom Davenport

To meet the Triple Aim, providers must compete not only on analytics but on predictive analytics. The overarching driver of adoption is the need to act proactively rather than reactively. For example, population and risk management rely heavily on the ability to predict which genetic, behavioral and socioeconomic markers

are likely to trigger which chronic conditions; which consumers are likely to respond to various programs; and what interventions work best to prevent or manage chronic conditions for these patients. These systems are helping both payers and providers to move from fee-for-service to risk-based contracting.

Continuous Capacity Planning

Tools that enable continuous capacity planning are one of the most recent applications of predictive analytics to emerge. Frontline clinical leaders need information on how their current patients' needs should be met over the next few hours and days. They also need smart tools to help them set capacity, recruit staff and deploy resources without incurring unnecessary costs. Resource deployment is a complex problem that is typically addressed by inconsistent, department-centric efforts. Without the ability to predict patient demand, most health organizations have no choice but to call in staff at the last minute, which increases labor hour costs, since managers must rely on agency nurses or overtime. Similarly, during low-volume periods, the only way to decrease costs is to send staff home, which can involve short-notice pay and affect staff satisfaction and retention.



Essential Components

An effective capacity planning tool should be able to forecast arrivals and flow over multiple time spans and translate this anticipated activity into facility and staffing requirements. Most tools on the market start by analyzing at least three years of historical admission, discharge and transfer (ADT) data to establish a baseline forecast. Capabilities vary widely, however. Some systems forecast future census, but not patient arrival times and flow during the patient's stay. Nor do they translate the forecast into a staffing plan. Some systems fine-tune the forecast

A continuous planning toolset should be accurate enough at each stage to help drive strategic, budgetary, scheduling, patient flow and staffing decisions, and flexible enough to enable operational and on-the-fly changes based on the complexities of the typical healthcare environment.

continuously in near real time, while others produce a single static forecast. Before you purchase, you should know if a system:

1. Manages multiple planning horizons: multi-year, annual, 6-8 week scheduling periods, weeks, days and hours
2. Continuously updates forecasts
3. Can forecast patient demand with a consistently high degree of absolute accuracy vs. a stated accuracy percentage that is aggregated over time
4. Forecasts volumes for door-to-door patient flow vs. preset intervals and departments
5. Converts forecasts automatically into capacity and staffing needs

Forecasting Methods

For more than a decade, many workforce management systems have included some degree of census forecasting ability. But while forecasting technology has vastly improved over time, most systems can still only produce very simplistic forecasts. The real benefits emerge only when forecasting is used over multiple time horizons (annual, quarterly, daily and real time), is integrated into the daily workflow of frontline managers upwards, and planning is fully aligned with process improvement efforts.

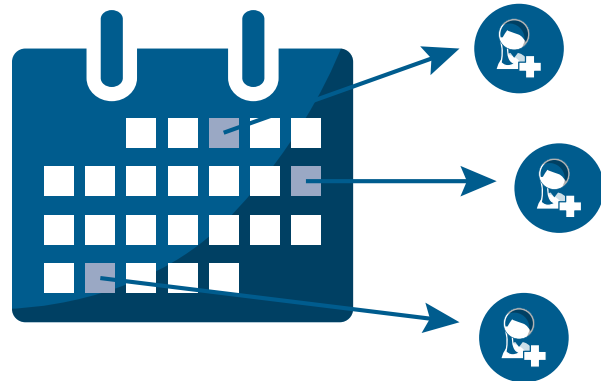
There are five key demand forecasting methods that together enable these benefits in a clinical environment:

1. Predictive Modeling. Predictive modeling is the use of historical data to determine factors that might cause something to happen. Commercial analytics platforms and business intelligence products lend themselves to the discovery of factors, but they cannot easily produce a patient demand forecast because they don't take into account patient volumes, arrival patterns and other important inputs.

2. Algorithmic Modeling. Averages and distributions of past activity can be combined in an algorithm to produce a forecast or prediction. For example, averaging the last 12 weeks of census levels by location, shift and day of week would be a simple algorithmic model that predicts future census levels. It might also include a trend factor if census has been rising or falling.

6. Supports user insight, input and adjustment as part of the planning process

A full-spectrum predictive analytics tool for capacity management should enable you to reliably forecast patient demand far enough in advance so you can align staff and other resources without incurring agency and overtime costs. It should be accurate enough at each stage to help drive strategic, budgetary, scheduling, patient flow and staffing decisions, and flexible enough to enable operational and on-the-fly changes based on the complexities of the typical healthcare environment. The forecasting methodology should be transparent to all stakeholders so that they trust the data and resulting projections. It should enhance the skill of clinical leaders by pointing out needed actions in advance and expanding the options available to solve operational problems.



Understanding the accuracy of your model is essential, especially if you are going to use it to allocate resources. For starters, getting a frontline care leader to trust a forecast is key to its use and impact. With statistics, it is easy to show a high level of accuracy. Make sure that accuracy is based on an absolute daily comparison of the forecast to what actually happened. Some accuracy analyses compare the average monthly census of the model to the average actual census. This aggregation trick can make a model appear much more accurate than it will be in practice.

3. Pattern Identification. Pattern identification is one of the building blocks of basic forecasts. Identifying trends and seasonal or day-of-week patterns in historical data is a proven forecasting technique. Likewise, removing anomalies from past data and flattening trends in areas that are reaching capacity is all part of the use of pattern identification in forecasting.

Once patterns are identified, multiple patterns can be combined into a forecast. This type of base forecast does not include things that have not occurred, and it might need further refinement to modify past activity. Flu season is a good example: the impact of last year's flu season is imbedded in the pattern, but a recent CDC report might indicate the need to intensify or moderate next year's flu season.



When leaders see how their input has been incorporated into the forecast, it helps to engender trust among the people who will use it to guide daily decisions.

4. Scenario Modeling. Once a base forecast has been developed, it must be refined by operational staff who can identify upcoming changes that will impact the forecast. Perhaps a group of surgeons is retiring, or a process improvement initiative is starting to shorten length of stay. Scenario modeling enables a planner to make such adjustments. In order to make the forecast more representative of what will happen, global adjustments can be made to variables such as inpatient/outpatient mix or length of stay. Other adjustments might only affect a date range, flu season or large community event.

5. Simulation. By combining arrival data (either using a statistical distribution or a known event) with predictive analytics about expected patient movement, a simulation can be developed to model and assess the impact of various scenarios. Care pathways can be combined with case booking or block scheduling data to create a simulation across departments. If the current OR schedule happens, what will be the bed and staffing needs on inpatient units? If a group of surgeons is hired or leaves, what will the impact be?

Simulations can also be used to test the impact of proposed changes before they are made. For example: a simulation of an OR schedule can forecast the activity on outpatient holding units, ICUs and acute care units. By making changes to the sequencing of OR cases, the simulation can look for sequencing that positively impacts census, perhaps by smoothing variation.

How to Compare Approaches

Now that you have a better understanding of various forecasting methods and their application to continuous capacity planning for healthcare, you are better equipped to evaluate your various options. As with any emerging technology, however, the same buzzwords may be used to describe vastly different capabilities, making it difficult for you to compare and evaluate options. This section compares various approaches and the limitations you may encounter when using a system that lacks one or more of the essential forecasting capabilities.

Aggregate vs. Patient-Level Demand

Some vendors “predict” patient-specific staffing needs a few days out, using a moving average patient census for the past several weeks. Staffing is calculated using tables set up in the staffing system, which may include acuity data. These systems are often integrated with a patient assignment module, and sometimes care paths tied to length of stay are used.

This approach is more accurately described as care planning, which may be tied to a rudimentary staffing schedule for the next few days. While invaluable for ensuring individual patients receive the right care and for

helping to predict discharge dates, care planning has little in common with operational predictive analytics, and it does not manage staffing costs proactively.

Static vs. Dynamic

Most forecasting tools on the market start by analyzing at least three years of historical census data to establish a baseline patient census forecast. They use a one-shot algorithm to predict census and a lookup table or care-hour-per-patient standard to calculate staffing requirements. As outlined above, a truly dynamic forecasting system uses multiple methods, including pattern identification, modeling, simulation and optimization. A trained planner carefully monitors trends and acts as a conduit between senior and frontline leaders. Feedback from operational and clinical experts is incorporated dynamically to refine the forecast based on their knowledge of expected or

You cannot impact operating costs without the ability to forecast across multiple planning horizons, continuously and as needed.

desired changes to service lines, payment models, competitors and other non-historical factors. Passive real-time adjustments are also used to keep the forecast continuously updated.

Short-Term vs. Full-Range Planning Horizons

Many vendors offer only short-term demand forecasting, a few days or weeks into the future. You cannot impact operating costs without the ability to forecast across multiple planning horizons, continuously and as needed. To cover the full capacity planning spectrum, you must be able to estimate your long-term strategic resource needs, establish an annual budget, publish schedules one month in advance and respond to change with continuous adjustments.

Strategic (3–5 years)

Strategic business plans typically include guidance around labor costs and capacity needs, including construction. Guidance is based on current trends, such as competitive and economic indicators that may drive patient volumes up or down. Many strategic plans end up collecting dust on executives' shelves after one review. The likelihood that the plan will drive strategy increases with the accuracy of the plan's information - including labor forecasts.

Budgetary (1–2 years)

Capacity planning often occurs wholly within the budgeting process for the upcoming year. Based on fiscal constraints, managers may be asked to cut last year's budget by a certain percentage, regardless of rising wages or competitive hiring pressures. Because most budgets are locked down upon approval, typically they are outdated as soon as they are created. Nonetheless, the budget represents the agreed-upon cost plan for each department. Managers are repeatedly asked to defend negative monthly variances (such as excess overtime) weeks after they occurred, and without access to tools that might help to avert these recurring crises.

Scheduling (1–4 months)

Whether they use pen and paper, spreadsheets or a software program, most organizations rely on the baseline ratios in the annual budget to calculate scheduling plans. Planning includes the review and comparison of actual activity to planned annual activity. The planning process should allow frontline leaders to address variation by adjusting staffing needs before schedules are created and posted. This will lower the need for "day of" adjustments and increase the use of straight time vs. bonus and overtime.

Operational (1-2 weeks)

Once a schedule is posted, many systems are programmed to simply wait until the beginning of each shift to react to and adjust capacity and staffing. The options that are available just before a shift starts are limited and costly. Continuous planning systems that accept real-time patient flow data can continuously adjust over the near term (3-7 days) and supply needed information to frontline leaders.

Continuous planning

The reality of a clinical environment is that unplanned things do occur. By employing a continuous planning process, adjustments on the day of can be reduced but never eliminated. Making these "one-off" changes part of the continuous planning process ensures they are rolled into future planning periods. Culturally, your organization might believe that a certain event happens so infrequently as to be inconsequential, or that the event recurs regularly but cannot be anticipated. Making these "one-off" changes part of the continuous planning process ensures that they are rolled into future planning periods and become part of the forecast.

A continuous planning toolset enables continuous learning – and increases the amount of time and options people have to prepare for what is no longer unexpected. A 3-5 day view into future demand and resource requirements gives leaders more time to consider a variety of less expensive options. For example, rather than adding last-minute staff to cover a spike in demand, you may be able to delay some elective surgeries or re-sequence some scheduled tests. If you still need to increase staff, you can do so based on straight time.

Top Down vs. Bottom Up

Because capacity planning is inextricably tied to staffing, too often changes are initiated by a top-down budgetary mandate to cut a certain number or percentage of FTEs without regard to patient demand or workload. Rarely is this approach sustainable. Planning must take a bottom-up approach to achieve the level of buy-in required for managers to rely on the system in their daily workflow and to preserve staff satisfaction. It must start by looking at how your organization aligns staff and resources with patient flow today – by arrival time, shift, hour and more. It must then consider various ways to hit that top-line metric, with some of the best ideas coming from frontline staff who know best where the bottlenecks and inefficiencies are.

OR Optimization

Most of the financial emphasis around capacity management focuses on reducing labor costs while optimizing patient flow. When the goal is to improve operating margins, then that approach covers only part of the equation. Even though smoothing the OR schedule is one of the best levers for smoothing patient census, most capacity planning tools do not address OR optimization, let alone its downstream impact. Too often, OR optimization is segmented off as a separate planning project, ensuring suboptimal results from the start.

The single most profitable area of any hospital is elective surgery. Some planning tools can predict the downstream impact of scheduled (or elective) surgeries. But a partial picture doesn't provide enough information to effectively plan for resources downstream. With enough historical data, as well as the right algorithms and ongoing expert input, unscheduled demand can also be predicted with surprising accuracy.

Other Considerations

In order to evaluate your needs, you must be able to assess a solution's ability to fully impact outcomes. In addition to forecasting methodologies, a number of other important factors must be considered to select the right solution.

Embedded vs. Agnostic

Some capacity planning tools are embedded in staffing or other systems. By definition, they can only be used with that system. These systems appeal especially in healthcare organizations that favor fewer vendors and fewer interfaces. The ideal capacity planning system should be vendor agnostic. It should have a slim footprint and rely on standard interfaces that work with multiple staffing, scheduling, patient flow and analytics systems. A capacity planning tool embedded in a proprietary staff scheduling system may be able to solve a limited number of needs related to that staff scheduling system's capabilities, but it cannot work with a patient flow system, for example, to ensure staff are aligned with other resources. Because embedded tools are typically enhancements to their parent systems, their costs often reflect their marginal utility.

To achieve sizable, sustainable gains, you must be able to predict scheduled and emergent demand by procedure, so that you both optimize the OR schedule and align resources in affected units.

To achieve sizable, sustainable gains, you must be able to predict scheduled and emergent demand by procedure, so that you both optimize the OR schedule and align resources in affected units. An OR optimizing tool should enable planners to build multiple what-if scenarios that illustrate the daily impact of proposed changes on other departments and units before the changes are made. For example, suppose that you can show a particular group of surgeons that by moving all of their procedures to a different day their patients will experience better flow in the perioperative suite and better care for the remainder of their stay. Chances are you will gain rapid agreement.

Granularity

Outside of consulting engagements, capacity planning is typically reviewed monthly or quarterly by a financial analyst using retrospective data, perhaps aggregated by day of the week. That may be adequate to present a snapshot of past performance, but it's hardly actionable. Financial analytics tools cannot prospectively model daily, let alone hourly, activity based on when patients are expected to arrive. Nor can they model the impact of a change in one department on other departments. Any frontline manager will tell you that this is the degree of granularity needed to resolve potential issues before they happen.

One of the keys to improving capacity planning and staffing is breaking through the shift mentality and focusing instead on patient arrivals and flow. Rather than reporting activity at the shift level, the system should provide nurse managers with real-time patient arrivals by the hour, with adjustments to the staffing plan already incorporated.

Transparency

Trusting the forecast is one of the biggest hurdles many organizations face when moving to a proactive capacity management approach. There's only one way to win trust:

by ensuring the forecast is consistently accurate and demonstrating that staffing to demand improves both staff and patient satisfaction. A good system provides the tools required to monitor progress against the forecast and adjust as necessary to stay on course. It enables you to quickly see where you will be understaffed or overbedded—and plan accordingly. As staff see the forecast tracking to actual demand and begin to experience a less chaotic environment, a proactive culture can take hold.

Consultant Driven vs. Vendor Supported

Health organizations have long turned to outside consultants for help with capacity planning. While approaches vary, many consultants provide patient demand forecasts and staffing schedules as part of their engagement. Relying on consultants who maintain control of their proprietary methodologies can be costly over time. More importantly, it can foster ongoing dependence without creating internal capabilities to manage something that is well within your organization's grasp and essential to your operations.

Still, for a continuous planning toolset to have an impact, it cannot simply be installed and turned on. It must be carefully integrated with your organization's routine workflows to enable the data-driven changes that will

Advisory services should help ensure that your designated planner has mastered the toolset, that the forecasts are accurate, that staff trust the data and have bought into the new process, and that you are on a path to achieve expected results.

Conclusion

As with any emerging technology, the market for capacity forecasting is rapidly evolving. Entrants include staffing and scheduling vendors, niche consultants and cross-industry analytics firms. The definition of terminology differs from vendor to vendor, and the same terms tend to describe widely varying capabilities. Some vendors embed forecasting tools in their existing solutions at little to no cost. While these vendors may purport to offer comparable results, the ability of their tools to produce useful forecasts and desired outcomes varies. Buyers must be able to identify what full-spectrum proactive capacity planning looks like, clearly define what they

occur. Your implementation should include advisory services provided by highly trained experts with decades of experience in hospital operations. They should work with you to define your desired outcomes and translate them into objectives and metrics that will be used to measure your results. These services should help ensure that your designated planner has mastered the toolset, that the forecasts are accurate, that staff trust the data and have bought into the new process, and that you are on a path to achieve expected results. While there are similarities between a traditional consulting engagement and the services provided with a continuous planning toolset, there is one big difference: at the end of the engagement, you possess the tools and the knowledge to be successful.

Return on Investment

Full-spectrum capacity planning systems that are



integrated with the frontline manager's daily workflow and coordinated with workforce management and patient flow systems can produce significant returns on investment. Reductions in the average length of stay in excess of 15% have been documented. Reductions in care hours per patient day of up to 10% have also been realized. The cost per labor hour has also been reduced through the control of overtime, bonus pay, and agency use. Qualitative impacts to the patient experience, quality, safety and professional satisfaction are commonly reported. These benefits largely stem from the proactive culture that arises from the availability of transparent, trusted forecasts. If your system is just another dashboard display in the C suite, your efforts might not have a similar impact.

need such a system to do, and challenge vendors' claims regarding what their tools can deliver what it can deliver.

The important thing to remember is that all forecasting tools are not created equal. Do your homework, be clear on what you need to accomplish, know which questions to ask and what to listen for in the answers.

¹“Risk and Reward: Positioning Your Health System to Deliver Value in a Transforming Health Care Marketplace.” The Advisory Board Company interviews and analysis, 2012.

²“The Cost of Caring,” American Hospital Association, March 2010

³ www.ihl.org/Engage/Initiatives/TripleAim/, accessed Oct. 15, 2014.

McKesson Capacity Planner™ is a continuous planning tool that enables you to reliably forecast patient demand far enough in advance so you can align capacity and staffing without incurring last-minute expenses. How? A combination of historical ADT data, algorithmic forecasting, clinical expert input and OR modeling results in a patient demand forecast with an absolute daily accuracy of up to 95% one month out and up to 97-98% one week out. It then leverages this forecast to predict capacity and staffing requirements. The result is more accurate capacity planning and staff scheduling, lower labor costs and less variability in labor costs.

For more information, visit www.mckesson.com/capacitymanagement

McKesson

5995 Windward Parkway
Alpharetta, GA 30005
www.mckesson.com

Copyright © 2014 McKesson Corporation and/or one of its subsidiaries. All rights reserved. All other product or company names are trademarks, service marks or registered trademarks of their respective companies.

3312-011/14



**BUSINESS
CARE
CONNECTIVITY**

MCKESSON