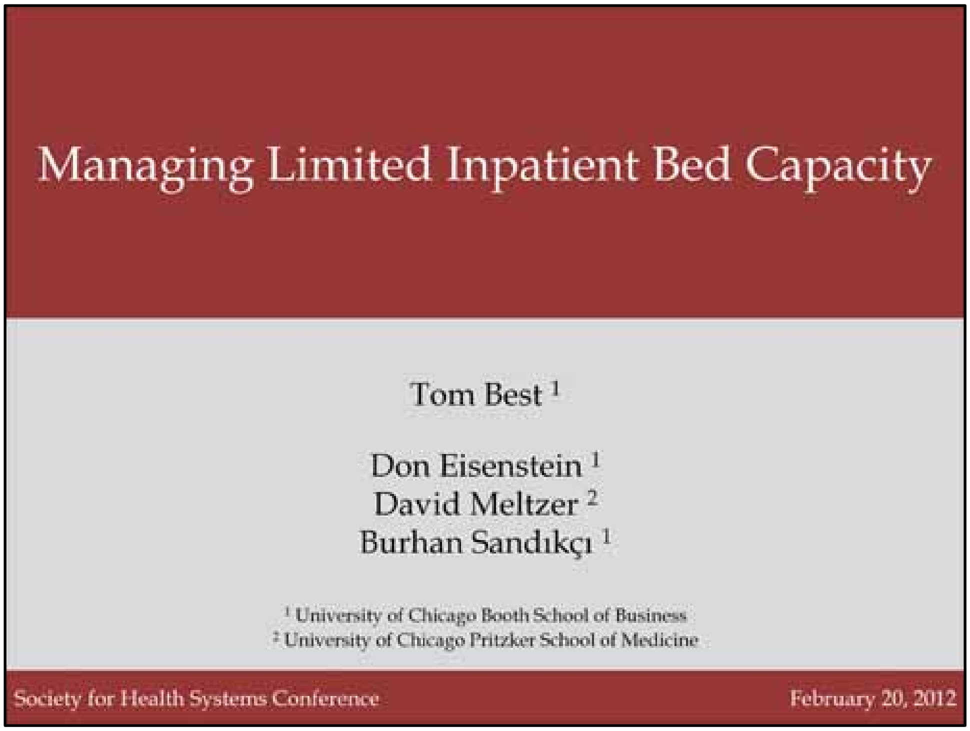

Tom Best is a PhD student in Operations Management at the University of Chicago Booth School of Business. In 2008 and 2009, Tom worked as a health systems engineer in clinical and operations quality at the University of Chicago Medical Center, a 596 bed academic medical center on the South side. He previously received a Bachelors (2006) and a Masters (2007) of Science in Industrial Engineering from the University of Wisconsin-Madison. Tom is also a member of the SHS Communications committee.

Listed below Tom are his advisors and coauthors: faculty at the University of Chicago.

This presentation discusses a model that is intended to assist hospital administrators in managing their bed capacity. The purpose is not to provide results for a specific hospital, although such results are provided, but rather to demonstrate the usefulness of the model and the feasibility and importance of using it in practice.

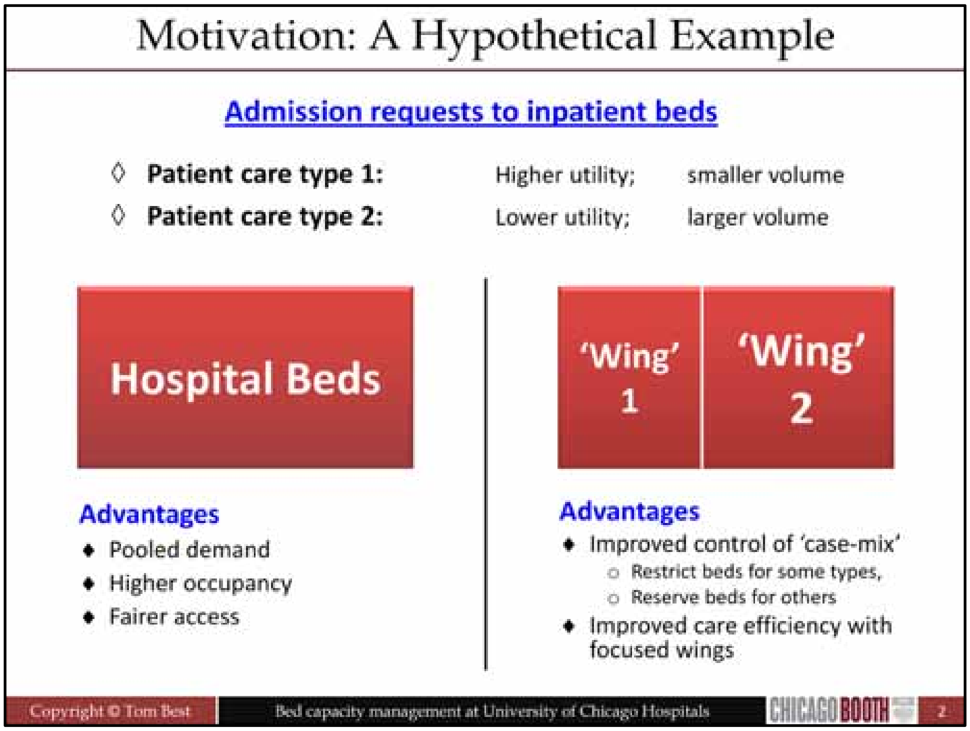

This hypothetical example illustrates the primary tradeoffs inherent to this research problem.

Consider two care types, or two different groups of clinical diagnoses (e.g. cardiology and neurology patients).

Suppose the combined demand volume of both care types is higher than the number of beds available.

On the left, both care types admit to one pooled set of beds.

On the right, care type 1 is assigned **only** to wing 1, care type 2 is assigned **only** to wing 2. Wing 1 is larger than necessary; in essence, the hospital is reserving beds for care type 1 and restricted beds for type 2. This leads to high occupancy and low access for type 2.

Which option is better: left or right? That is the question we will address with our research.

Motivation: The Real Case

- ◊ **The University of Chicago Hospitals (UCH)**
  - ♦ High demand for their inpatient beds, yet
  - ♦ Difficulty admitting a case mix that best supports their missions
    - First class patient care
    - Clinical research
    - Medical education

 Bed capacity management at University of Chicago Hospitals CHICAGO BOOTH 3

This example was motivated by a real clinical scenario.

Instead of 2 care types, consider about 25 care types, all of which have “faculty physicians” who aim to conduct outstanding patient care, research, and medical education in their discipline. Yet the demand for a few care types far surpasses the demand for other care types.

This is the situation at the University of Chicago Hospitals.

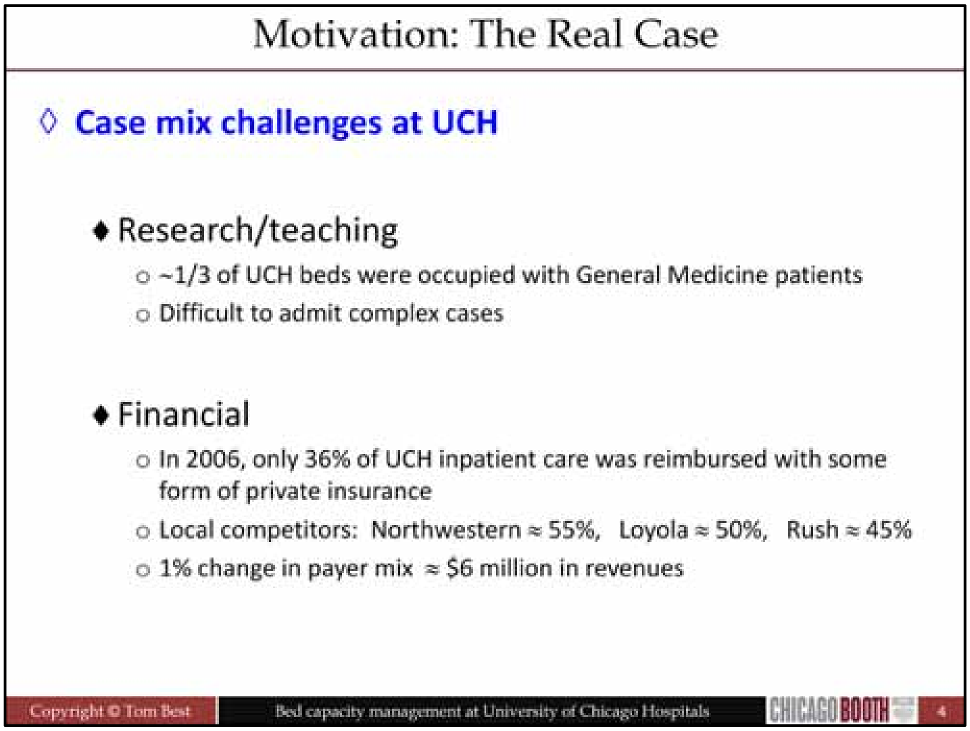

In more detail, here are the case mix challenges.

In terms of research and education, there are a number of local factors putting strain on the U of C.

- 7 South Side hospitals have closed since 1986 (CBS Chicago 2/15/11 article http://chicago.cbslocal.com/2011/02/15/ambulance-runs-end-at-provident-hospital/), the seventh: Oak Forrest (reference: http://www.suntimes.com/news/metro/4854787-418/patients-brace-for-oak-forest-hospitals-likely-closure.html from 8/4/11)

- Given U of C's strong reputation as a top provider of care on the South Side, this translates to an increase in local demand for care, care which does not always require the research- and education-oriented cost structure that U of C maintains.

A low private pay percentage in comparison to their local academic competitors made it even more difficult to continue to meet all three of their missions. Top research and education environments are difficult to maintain without the revenues achieved by your local competition.

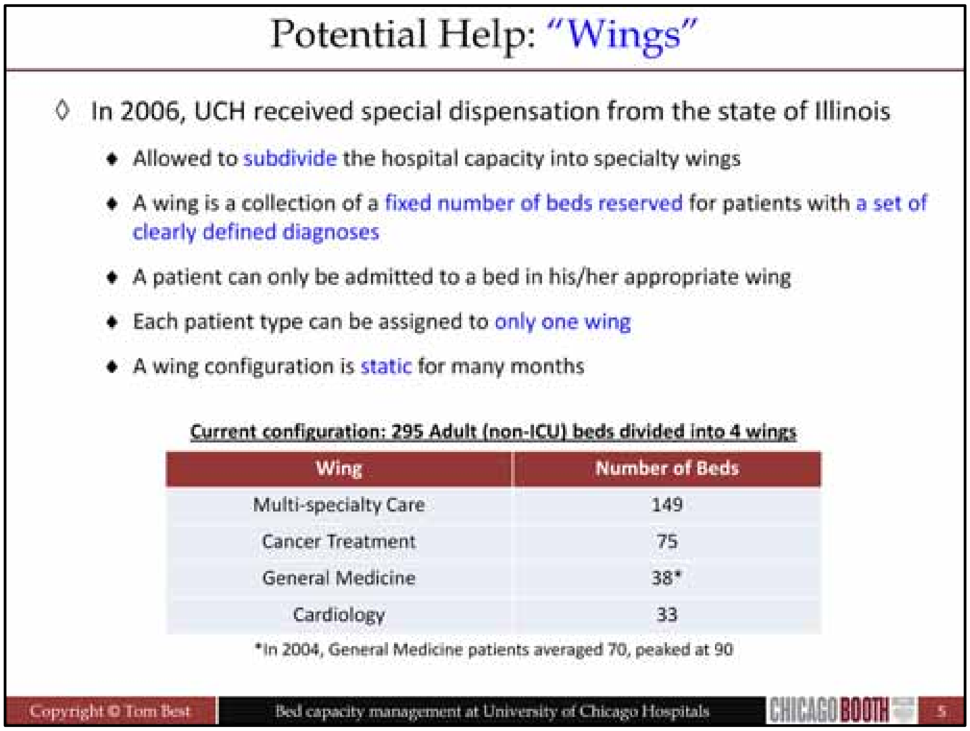

If a wing mismatch is identified, the inpatient must be transferred to their appropriate wing as soon as a bed is available in that wing.

The wing configurations are static due to the need for dispensation from the state.

We see from the asterisk that General Medicine is likely one of the groups whose bed capacity has been restricted.

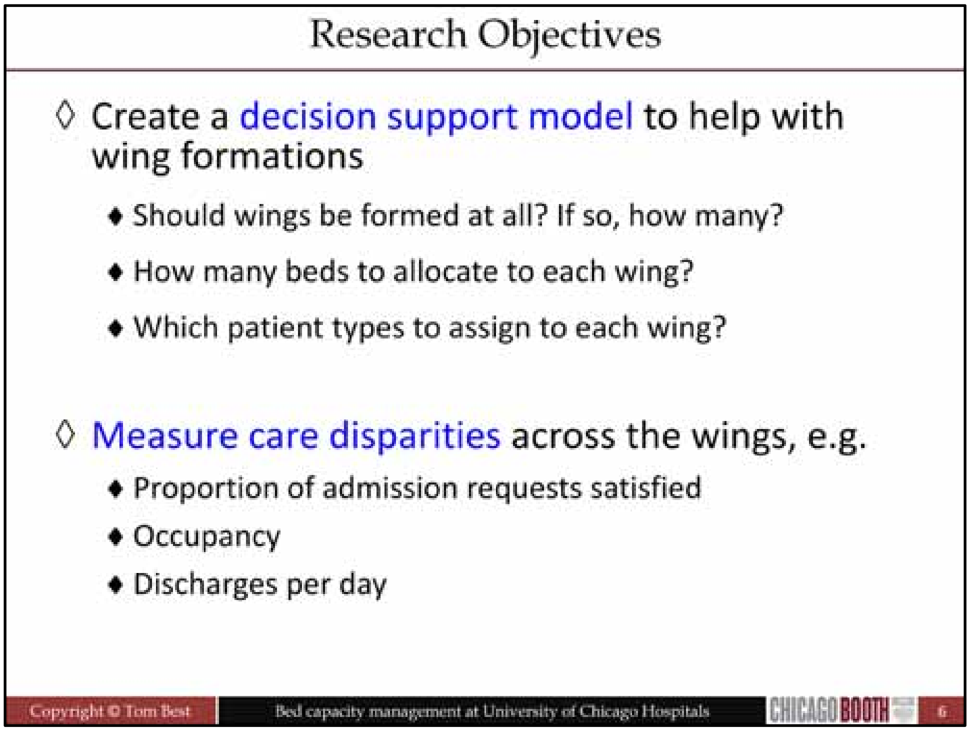

Our perspective is that of the hospital administrator. Both of these objectives have equal importance from this perspective.

The model discussed from here forward will be an optimization model, not a simulation as may be more familiar to some. The value of using this model is that it gives the best solution under a set of inputs. We will formulate a model of the problem faced, and find the best solution, subject to the modeling assumptions and the inputs (e.g. the patient types, admission request demand, length of stay, utility). We will review in more detail later.

(Other care disparities, such as expected waiting time, are discussed during the actual presentation.)

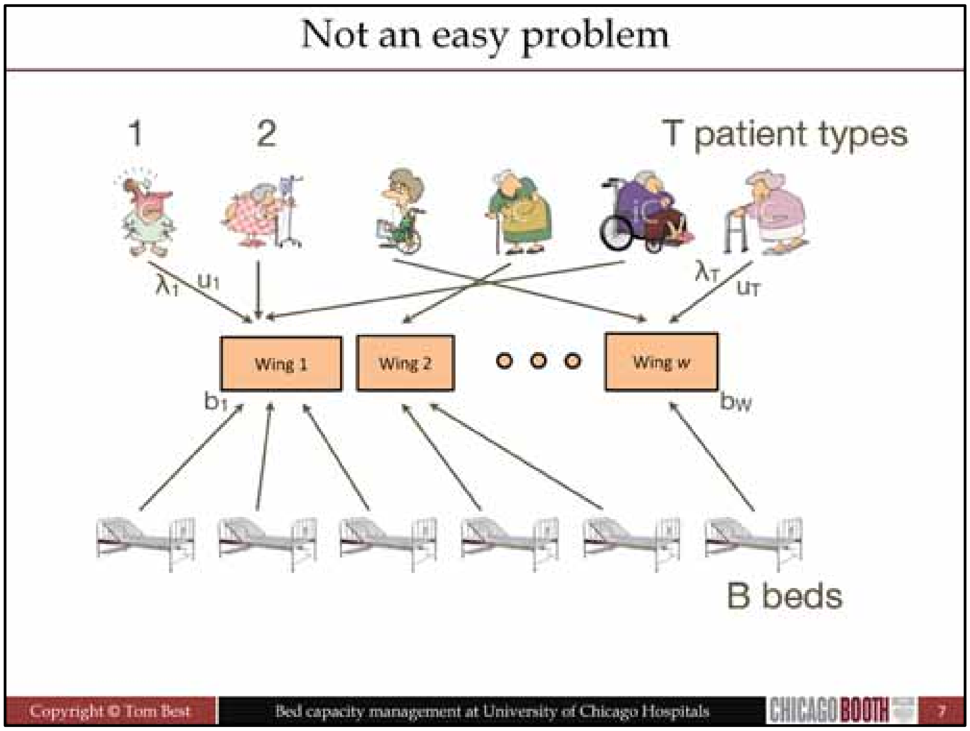

In order to build a useful model, we must decide how to group T care types (e.g. T = 24) into w wings, where w can be from 1 to T. The model will decide how many wings to create, and also how to allocate beds to each wing.

Each care type has a different demand, utility for the hospital, and length of stay.

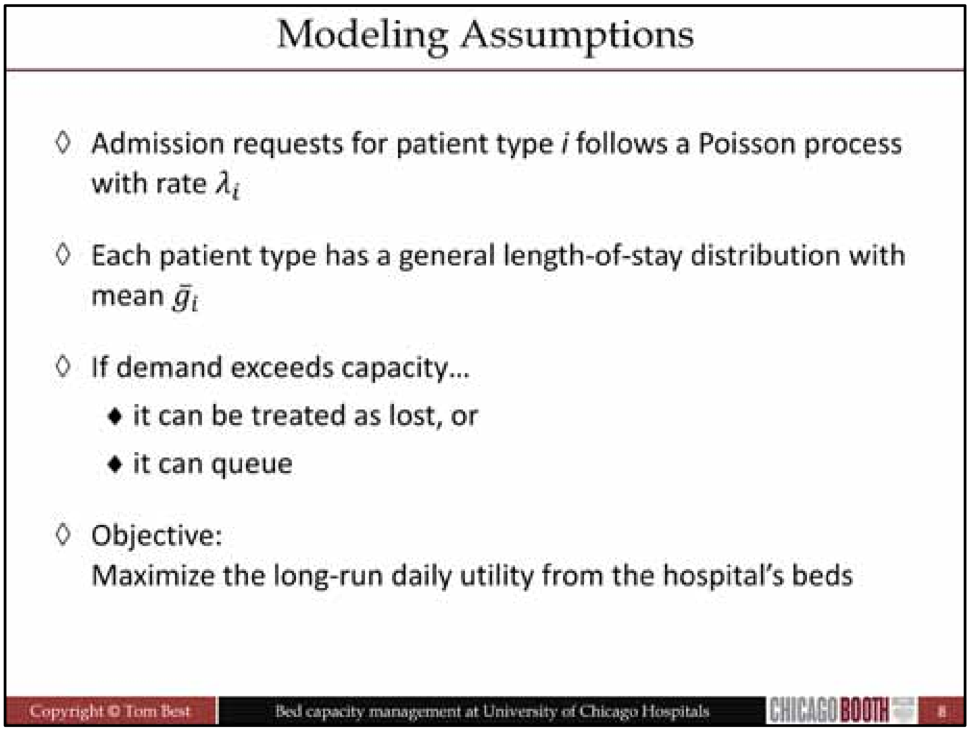

A Poisson process means the time between new admission requests is random, following an exponential distribution.

The average or mean LOS is for any distribution you'd like (although if it's important that admission requests want to be accepted for some amount of time (e.g. a day, week, or more) when the beds are full, we will assume length of stay is exponentially distributed as well.

The utility is for the hospital (one example of utility could be reimbursement for a patient's stay, but we will discuss another that we use in the coming slides).

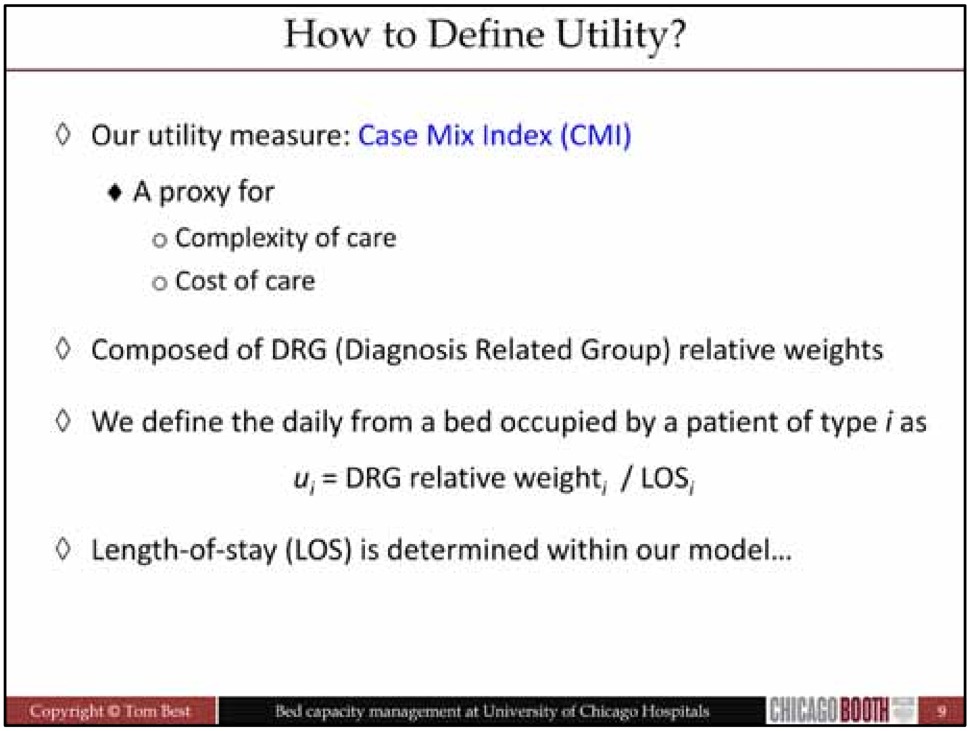

Using data from the U of C, we estimate the average DRG Relative Weight for inpatient episodes within a given care type, as well as some nominal length of stay for the wing.

While it is solving, our model then adjusts length of stay endogeneously, i.e. depending on the wing's structure. We discuss this in the coming slides.

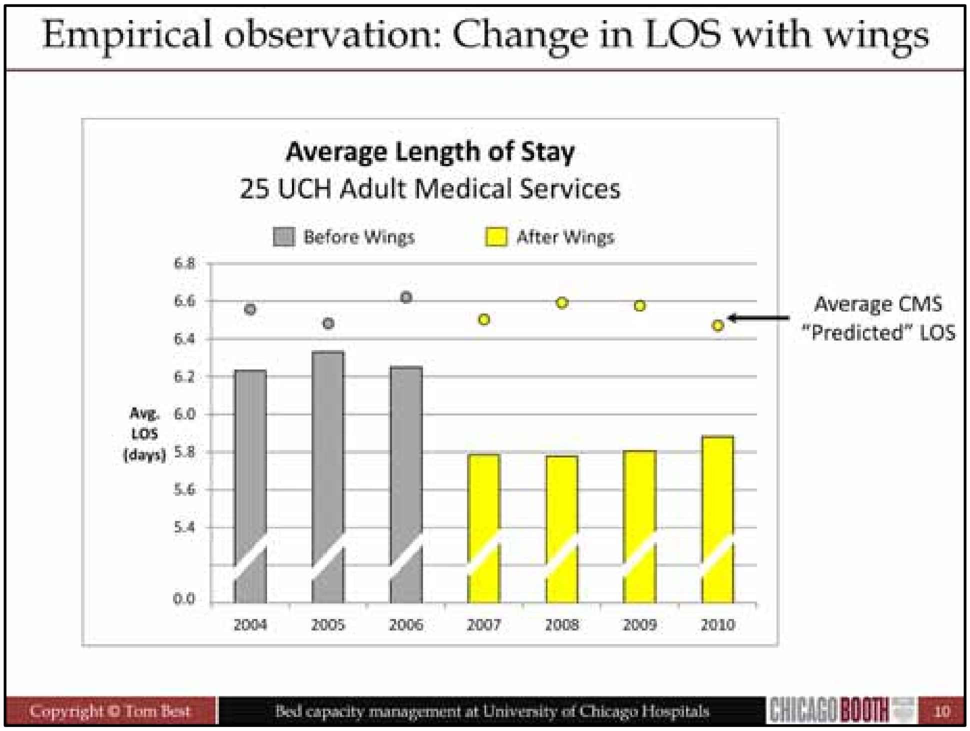

Here's something we find in our data: there was a decrease in overall average LOS at UCH associated with the introduction of wings which was robust to changes in care population and care technologies over time.

Why did this drop occur? Was it truly associated with the introduction of wings, as suggested by this graph?

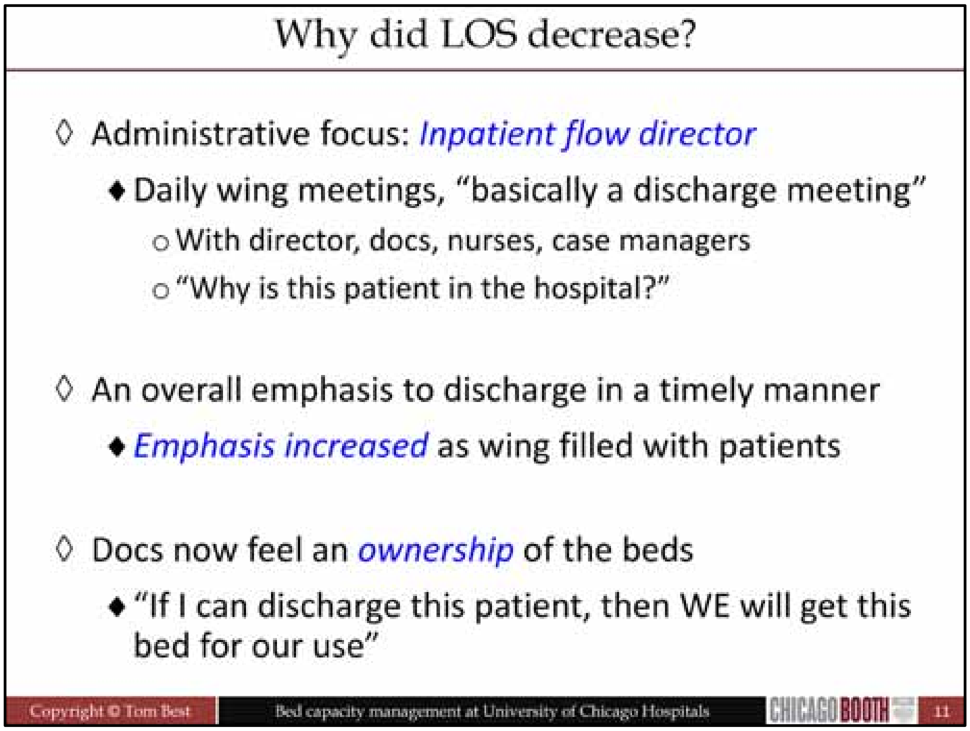

To answer the previous slide's question, we talked with a long-time director of one of the wings. Shown on this slide is what this person had to say.

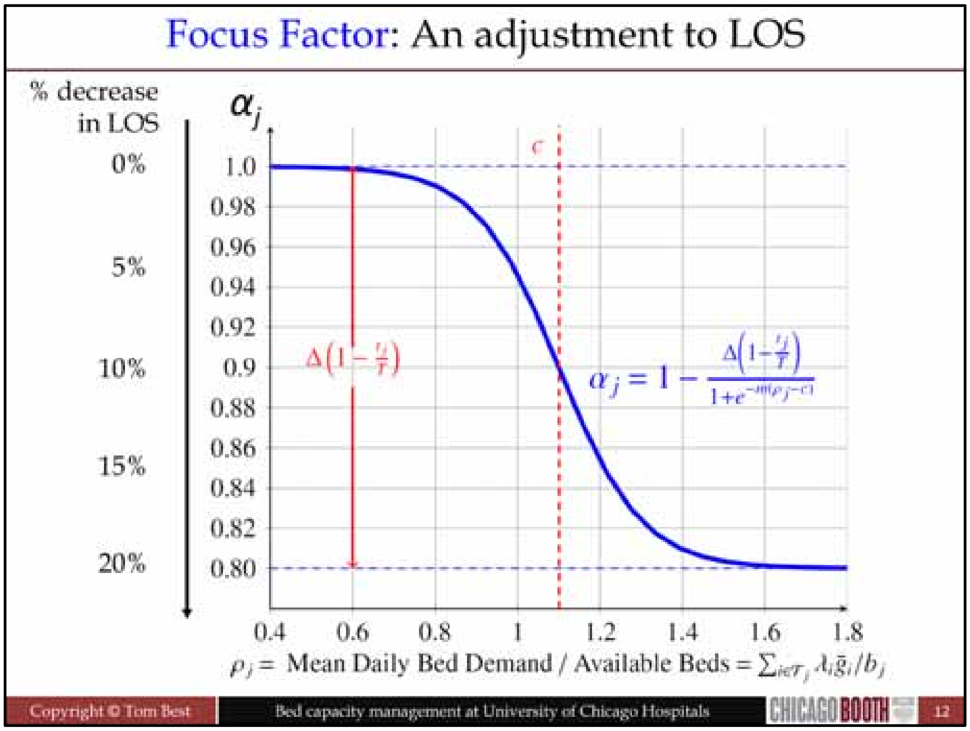

Here's how our model adjusts length of stay (LOS) depending on the structure of the wing.

To summarize the detailed picture, we use this "focus factor" in the following way in our model:

- When the average number of beds demanded (which is at least as high as the number of beds occupied) is higher than the available beds, average LOS for the wing will be lower, to reflect the emphasis on discharge described by the inpatient flow director (and seen in other literature) when the wing's beds are in high-demand.

- When there are only a few care types in the wing, the decrease in length of stay is more significant when there are many types in the wing. In fact, when all types are in one wing, the focus factor is null (there is no reduction in LOS due to the creation of wings, because no care types are allocated separate sets of beds).

This picture is the "focus factor" we thought best represents the LOS data we've analyzed, and the discussions we had with the inpatient flow director. The model's relationship between LOS and wing structure is much more than flexible, and the model user can tweak a number of parameters to decide what is best given their hospital's "particulars".

Note: we assume these reductions in length of stay are associated with reductions in the non-value-added time during a patient's stay, NOT by "cutting corners" and discharging a patient before they are ready. A maximum decrease of 20% as shown here is extreme – in reality, we estimate these reductions to be as high as 10%, and more often 5% (1 in 20 days of care).

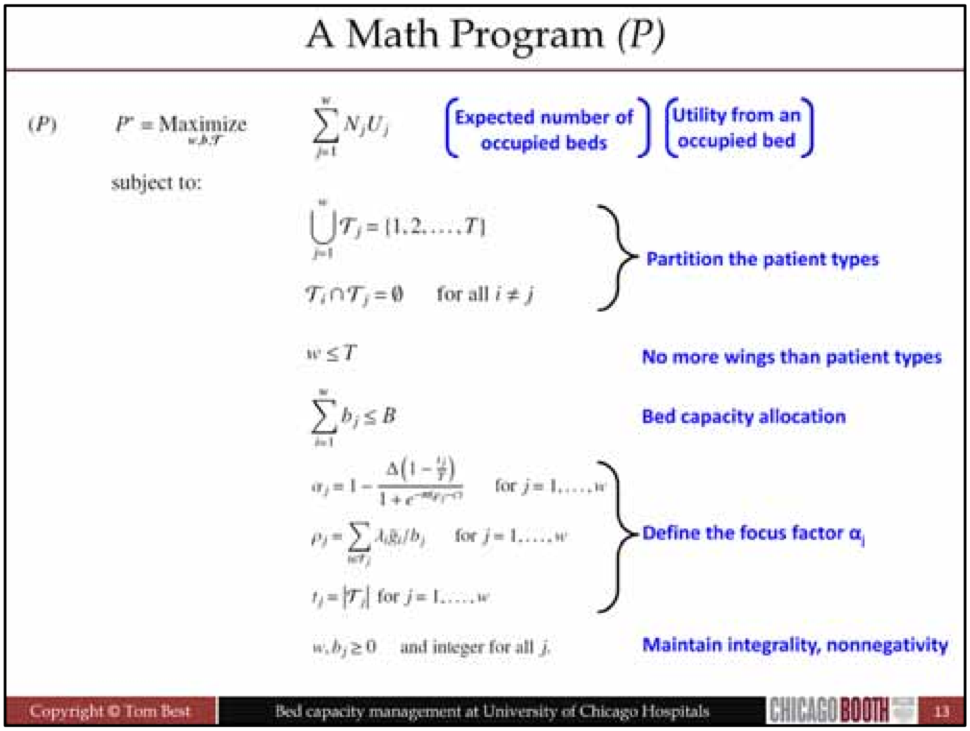

Now, for those interested, we are ready to show our overall math model.

We summarize each equation in words, using the right hand side of the slide.

The model takes as inputs:

1) A defined set of inpatient care types (e.g. cardiology, general medicine, orthopedics, etc),
2) A measure of utility for the hospital for a given inpatient encounter (e.g. DRG relative weight, reimbursement from the insurer, etc.)
3) For each type
    1) An estimated average number of admission requests per day that would be made without knowledge of the current idle/available resources at the hospital (beds, nurses, etc) (these requests are NOT necessarily accepted)
    2) An estimated nominal length of stay (the average length of stay before wings have been created)
    3) An estimated average utility per encounter
4) A total number of beds to be allocated (not necessarily the total number in the hospital), ideally that are all considered “equal” (in that any bed can be used for any care type). This last requirement is not entirely necessary.

Given these inputs, the model chooses the number of wings, the types in each wing, and the number of beds in wing, so as to maximize total long-run average daily utility for the hospital.

For example, the objective function is the sum of the utility contributions from all the w wings.

The model must satisfy some requirements, or constraints, when evaluating a solution:

- Each care type can be assigned to one and only one wing.
- The total allocated beds can’t be more than the number B available.
- The rest of the constraints are just to define the focus factor, and to maintain an integer number of beds and wings.

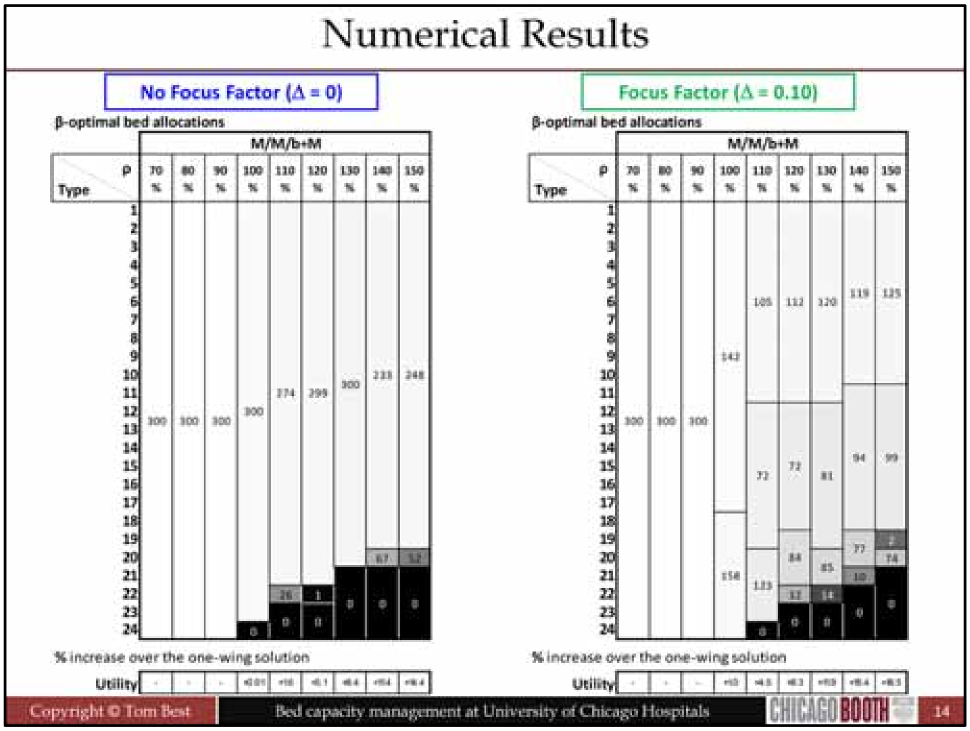

Here are some solutions recommended by our model for various sets of inputs. Consider this slide a "sensitivity analysis": how sensitive our the model solutions to the model inputs?

Each column represents an optimal solution under a set of inputs.

The set of columns on the left are with no focus factor, that is, no decrease in length of stay due to the formation of multiple wings.

The set on the right has a maximum decrease in length of stay of 10%, which only occurs for a wing with 1 care type and very few beds in comparison to daily bed-demand. Otherwise, the decrease in length of stay is actually much less for a given wing (say only 5%).

To describe one solution, consider the 100% column in the set on the right. In this scenario, there are 300 beds to be distributed amongst 24 care types, which are ordered from highest daily utility to lowest daily utility. Overall average bed-demand is 100% of bed capacity (300 beds demanded on a given day). An average admission request is willing to wait 3 days. If not accepted/admitted after 3 days, the person goes and seeks care elsewhere. The model assigns 142 beds to the 17 care types with the highest daily utility, and the remaining beds to the 7 types with the lowest daily utility. This results in a 1.0% increase in daily utility compared to the solution which assigns all the types (and all the beds) to one wing.

The shading for each wing indicates the service disparities across the wings. A darker wing receives poorer service from the patient's perspective: higher occupancy than lightly shaded wings, and (in general) longer waits for admission.

We see two main messages.

- When demand is at least as high as capacity, the model assigns less capacity to the lower-utility types, and this improves utility. Otherwise, it's best to just group all care types into one wing, because we see more patients that way.
- Considering focus is important. More wings are formed in general, and this results in greater improvements in utility, but more significant disparities in patient service.

The model solves one input scenario in about 2 minutes on a personal computer.

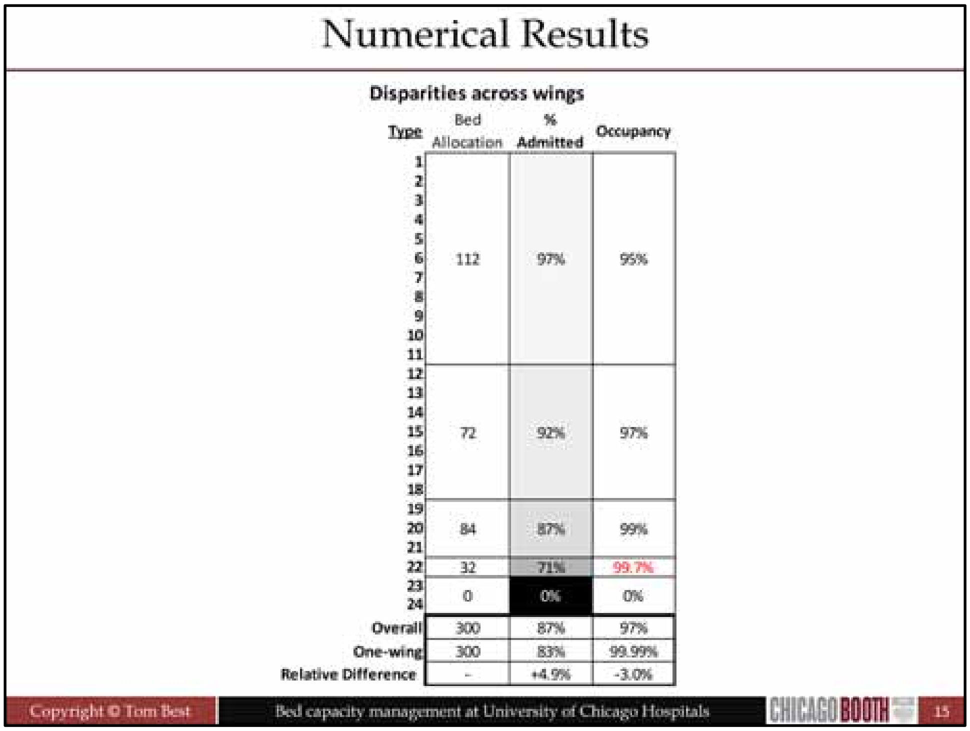

Type	Bed Allocation	% Admitted	Occupancy
1–11	112	97%	95%
12–18	72	92%	97%
19–21	84	87%	99%
22	32	71%	99.7%
23–24	0	0%	0%
Overall	300	87%	97%
One-wing	300	83%	99.99%
Relative Difference	-	+4.9%	-3.0%

Here are the disparities shown with numbers in addition to shading. We see that admissions and occupancy change significantly across the wings: the highest utility types are given preference.

Technical details of the shown model input scenario:

- Patients wait for care for an average of 3 days,
- Average daily bed-demand is 120% of capacity,
- The max decrease in LOS is 10%,
- There are 300 beds available
- The patient types are ordered from highest (1) to lowest (24) daily utility.

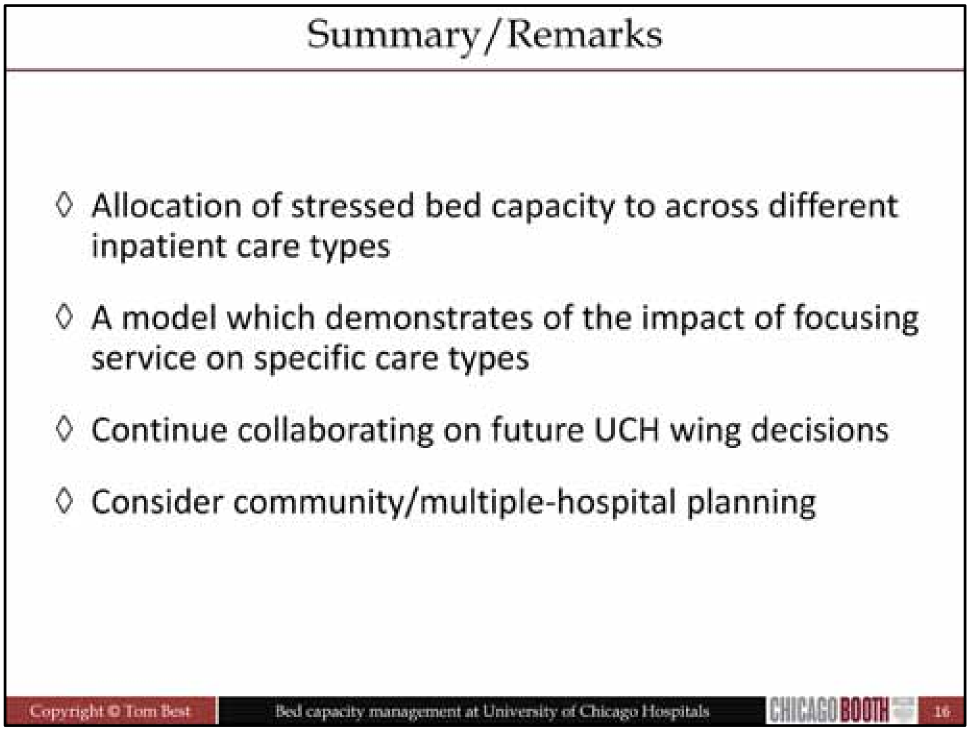

In conclusion, this presentation has discussed a model intended to assist hospital administrators in managing their bed capacity. We have seen that the model is useful for assisting with this difficult decision, is feasible to use in practice, and is important to use, since key implications of the recommended solution are provided as part of the model outputs, and may be very influential toward the success of a hospital in achieving its missions.

The largest barrier to utilization of this model (as is true for many models) is defining the correct model inputs.

It requires significant, complicated data analysis and at least three critical decisions before using the model; in particular,

- how should the care types be defined,
- how should true demand be assessed,
- and what should be the utility measure for a given patient care episode.

For that reason, we are developing a supplementary project to describe defining these inputs for a given hospital in more detail.

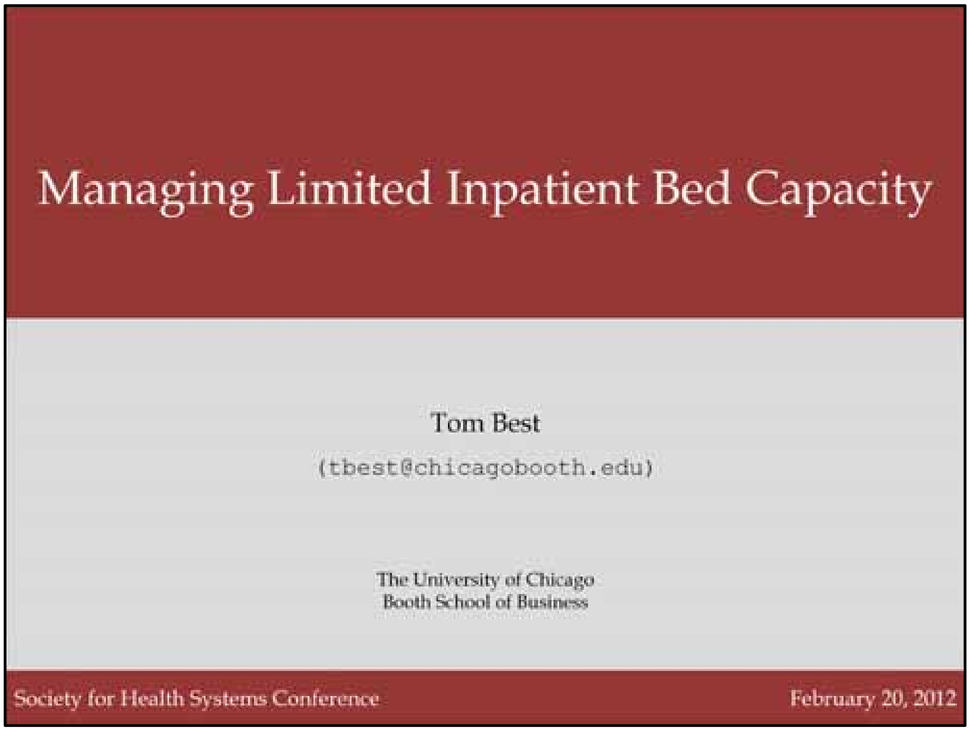

Please feel free to contact me with any questions or comments.

Tom Best
tbest@chicagobooth.edu