

Applications of Operations Research in Real-Life Scenarios: Hospital Resource Allocation (Linear Programming and Goal Programming)

Authors name: **Manan Beria, Naman Shah, Nandini Moorthy**

Course & College: Bsc. Finance, NMIMS Deemed to be University



<https://doi.org/10.55041/ijssmt.v2i10.037>

Cite this Article: Beria, M., Shah, N. & Moorthy, N. (2026). Applications of Operations Research in Real-Life Scenarios: Hospital Resource Allocation (Linear Programming and Goal Programming). International Journal of Science, Strategic Management and Technology, 02(10), 1-9.
<https://doi.org/10.55041/ijssmt.v2i10.037>

License:  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

Abstract

A hospital has a static number of ward beds, intensive care unit (ICU) beds, operating theatre (OT) hours, nurses and doctors and a weekly problem of allocating these resources between various departments and patients groups. This pressure is particularly strong in India, where there are roughly 1.6 hospital beds and 0.7 doctors available per 1,000 people, which is significantly lower compared to most of the larger economies in the world. This report focuses on the use of Operations Research (OR) in making these decisions and specifically on linear programming (LP). In depth analysis of two journal articles. The first, published in the European Journal of Operational Research by Blake and Carter (2002), is a linear goal programming model for determining the case mix of a Toronto teaching hospital following an 18% reduction in funding. In Naval Research Logistics, Mehrotra et al. (2020) develop a stochastic mixed-integer linear program to help allocate and share the stockpile of ventilators in the United States during COVID-19 between the states. For each paper we describe the real-life problem, the model, the solution method and its strengths and weaknesses. We then formulate a general LP for the problem of resource allocation at a hospital and solve a numerical example of our own, using Python. The illustration defines the binding resources and prices them with shadow prices, examines the sensitivity of the ICU capacity, and demonstrates how a 10% reduction in the nursing hours makes it necessary to make a decision between financial balance and service volume through goal programming.

1. Introduction to OR and the Problem Type


1A. What is Operations Research?

Operations Research is the use of mathematical models and analysis to decide how scarce resources should be used. The steps of a typical study are the same: The actual problem is formulated; The problem is written as a model including its decision variables, objective function (to maximise or minimise), and constraints - descriptions of the limits; The model is solved, either exactly, or using a heuristic; The answer is verified by comparing it to the actual situation; Implementation of the answer. The most popular of these techniques is linear programming, which was developed late in the 1940s by Dantzig. An LP's objective and its constraints are linear, and can be solved in seconds with a simplex method or an interior-point method when we have many thousands of variables. LP also provides more than just the answer—the dual values (shadow prices) indicate the value of one additional unit of each resource to a manager.

1B. Hospital resource allocation as an OR problem

The hospital resource allocation model is a natural LP problem. The decision variables are the number of patients of each type that the hospital will treat, or the quantity of some resource that it provides to each unit. Patients of each type utilise a fixed quantity of each resource (bed-days, ICU days, theatre time, nursing hours) and each resource has a fixed capacity, and the hospital wishes to maximise some measure of value (such as patients treated, benefit to health, or contribution to the finances). The same structure as the classic product mix problem in manufacturing, but patient groups take the place of products. In hospital planning, this is referred to as case mix planning and places decisions at the strategic and tactical level of hospital planning, rather than at the day-to-day level of planning.

It's not just an academic issue. The health-system capacity of each country is compared in figure 1. India's health infrastructure is severely underdeveloped, with the availability of just one hospital bed per 1,000 persons compared to over 12 hospital beds per 1,000 persons in Japan and South Korea and 2.68 beds per 1,000 persons in the United States, while there are just 0.72 physicians per 1,000 persons against 3 to 4.5 physicians per 1,000 persons in most of the high-income countries. If the space is so constrained, it's as much a matter of how it's divided as it is of how much there is.



Source: World Bank, World Development Indicators (SH.MED.BEDS.ZS, SH.MED.PHYS.ZS), latest year available (2020 to 2023).

Figure 1. Hospital beds and physicians per 1,000 people, India and selected countries

The models in this report differ in three ways from hospital allocation of resources because they are different from a factory product mix: The first is that there are more than one decision maker: doctors are concerned with their practice and income, the administrators with the budget, and the public with access. Several targets are often being dealt with in one problem - this is what is called goal programming (Charnes and Cooper, 1961), a variant of LP. Second, minimum service levels are imposed: a public or teaching hospital is not allowed to simply cease care of low-margin patients. Third, demand is uncertain, and in a crisis such as a pandemic the uncertainty becomes the main issue, which calls for stochastic programming. These are the angles picked – in Paper A, the multi-goal LP is within one hospital, in Paper B, the stochastic LP is a health system.

1C. Method and structure

Two journal articles were chosen for analysis, with the same headings used: real-life context, modelling, solution methodology and critical analysis. Analyses are based on published papers, their abstracts and bibliographic information and secondary literature on the papers. Papers are introduced in Section 2 and analysed in Sections 3 and 4, a general mathematical model is given in Section 5, this is then applied to a numerical example of our own in Section 6, and the approaches are reflected on in Section 7, before the conclusion in Section 8.

2. Selected Research Papers

Table 1. Articles analysed in this report

Paper	Reference	Domain	OR technique
A	Blake, J. T., & Carter, M. W. (2002). A goal programming approach to strategic resource allocation in acute care hospitals. <i>European Journal of Operational Research</i> , 140(3), 541-561.	Healthcare; hospital case mix and budget planning	Linear goal programming (two linked LP models)
B	Mehrotra, S., Rahimian, H., Barah, M., Luo, F., & Schantz, K. (2020). A model of supply-chain decisions for resource sharing with an application to ventilator allocation to combat COVID-19. <i>Naval Research Logistics</i> , 67(5), 303-320.	Healthcare; emergency medical supply allocation	Two-stage stochastic program solved as a mixed-integer linear program (MILP)

The two papers complement each other. Paper A allocates the resources of one hospital over a planning year under several conflicting goals. Paper B allocates a single critical resource across a whole country day by day under uncertain demand.

3. Analysis of Paper A: Goal Programming for Hospital Case Mix

3A. Real-life context

At the time of the study, Ontario hospitals in Canada were experiencing major reductions in public funding. The funding cuts for the teaching hospital Mt Sinai Hospital in Toronto were announced as a three-year drop of 18%. The hospital's surgical department had to determine which patients and in what numbers to see, to make up for the lesser revenue. This decision is the hospital's case mix. It had to be done without losing the surgeons, whose livelihoods and work modes rely on the numbers and nature of the cases they can be given. The field is healthcare management and public-sector budgeting. The problem has many patient types, many surgeons and several limited resources, all with linear relationships (cost per case, beds per case, etc. theatre time per case), which is what an LP can represent and solve.

3B. Modelling and problem formulation

Two linear goal programming models are given in the paper. The first one is when the case mix and volume for each physician are fixed. The second deals with the case mix selected and calculates the changes in physician practice (e.g., length of stay, tests used, etc.) that would be required to attain the cost targets. The main elements are:

- Objective: Minimise the weighted sum of unwanted deviation from a set of goals instead of maximise a single quantity.
- Goals and constraints: the hospital must break even (Case mix consistent with reduced budget); individual physician income should be close to a target; the bed, theatre, resource utilization must remain within availability; and the new case mix must not be too far removed from current practice.
- Decision variables: number of cases of each of the two types assigned to each physician (Model 1), and the change in resource use per case (Model 2); a pair of deviation variables (under- and over-achievement) for each goal.
- Type of problem: A multi-objective linear program in goal programming form, is a case mix (product-mix) problem in resource capacity.
- Assumptions: resource use is linear and fixed per case type; size of the population to be served by the planned volume of each case type is reasonable; goals and weights can be predetermined with the decision makers.

3C. Methodology

After including the deviation variables, a goal program becomes an ordinary LP, which can be solved using the simplex method in a typical LP solver. The contribution of the paper is the formulation and the way it is used rather than a new algorithm. Models were executed using data from hospital and physicians from the surgical division of Mount Sinai. They were validated with use: the hospital decision makers used the results to consider options and the models were used to identify the trade-offs between case mix and physician practice variables that were difficult to discern by hand.

3D. Findings, impact and analysis

The authors demonstrate how the models can be used by a hospital to determine its case mix and case costs as needed, without altering physician income significantly, and without requiring a great deal of change in practice. The method brings out the financial vs. professional motivations and puts them in black and white, allowing managers to discuss the numbers rather than opinions with doctors. One of its key tenets – that modest changes in case mix are more likely to be accepted by physicians – has been echoed later by Blake and Carter (2017) who point out that physicians will accept changes in case mix only when the changes are modest, which is one reason to model goals and deviations (and not attempt to force a pure profit-maximising case mix).

- Strengths: It can deal with multiple stakeholders simultaneously, it is based on standard, quick LP machinery, it generates trade-off curves that can be used for negotiation, and it has been tested on real hospital data, in a real context of a funding cut.
- Assumptions: Linear costs do not account for economies of scale and fixed costs for opening or closing services; the length of stay and demand are assumed; goal weights are subjective; emphasis is given to the surgical division of the hospital and not the entire hospital.

- Extensions: stochastic demand and length of stay; integer variables for discrete decisions, e.g. opening a ward; patient outcomes/waiting times as explicit goals; linking case mix planning with theatre scheduling, which was later tackled by the same hospital using integer programming.

4. Analysis of Paper B: Ventilator Allocation during COVID-19

4A. Real-life context

The United States was projected to face a shortage of mechanical ventilators in March 2020. The Federal Emergency Management Agency (FEMA) had 20,000 ventilators in its inventory and had then distributed over 8,100, with 4,400 units distributed to New York. Demand forecasts for the various states were different for the peak demand in April and May. A unit is wasted which goes to a state that already has reached its saturation point, while another state is lacking in capacity. Paper B focuses on where the central stockpile should go, and when states should return their idle ventilators to be used elsewhere. The field of work is healthcare logistics and emergency response. OR is appropriate since this is a large resource allocation over time with balance constraints, and since uncertainty of demand can be modeled in the form of scenarios within an LP.

4B. Modelling and problem formulation

The model plans for each US state on a day-by-day basis for 70 days (from 23 March to 31 May 2020). It is briefly summarized in Table 2.

Table 2. Main notation in Paper B

Symbol	Meaning
x_{nt}	Ventilators sent from the central agency to state n at the start of day t
z_{nt}	Ventilators returned by state n to the central agency on day t
y_{nt}	Ventilators in state n available for COVID-19 patients at the end of day t
s_t	Central (federal) inventory at the end of day t ; $s_0 = I = 20,000$
d_{nt}	Uncertain number of critical patients needing a ventilator (scenario-based)
Q_t	New production arriving on day t (100 per day, rising to 300 from 15 April)
γ_n	Share of a state's own stock reserved for non-COVID patients
τ_n	Share of usable stock a state is willing to share
ρ_n	Safety-stock multiplier a state keeps before releasing units (set to 1.5)


The goal is to minimize the expected total shortage, $\sum_t \sum_n E[(d_{nt} - y_{nt})^+]$ in ventilator-days. The constraints are: an inventory balance for each state (stock today = stock yesterday + received - returned), an inventory balance for the central stockpile (including production and returns), a limit on the number of units the agency can ship (its inventory), a sharing rule that a state returns units only from stock above its safety level $\rho_n d_{nt}$ and within its willingness to share τ_n , and initial conditions $y_{n0} = (1 - \gamma_n)Y_n$ and $s_0 = I$. Key assumptions include a single decision maker, decisions at the beginning of each day, a finite number of situations of demand that are known and have known probabilities, and no transport time and no loss of equipment.

4C. Methodology

The shortage term $(d - y)_t$, and the sharing rule, are non-linear, and are thus linearised by the introduction of an auxiliary shortage variable $ent_t = d_{nt} - y_{nt}$, $ent_t \geq 0$, and the replacement of the second sharing rule with a binary variable, along with big-M constraints. The stochastic program is then formulated in extensive form as one large mixed binary linear program, since there is a finite number of scenarios. The model was implemented in Python 3.7 and solved using the Gurobi 9.1 solver on a 14-core workstation, using a 1 hour time limit per run. For each of the four cases, between average and severe demand scenarios were considered from the IHME COVID-19 forecasts' confidence intervals (24 scenarios per case). To check the demand inputs, the authors compared the forecast with reported New York deaths (217.9 forecast and 237 reported on 29 March; and 262.2 forecast and 253 reported on 30 March)

4D. Findings, impact and analysis

The big issue is that the solution is largely state-dependent, in that the proportion of state supply dedicated to non-COVID patients plays a significant role. The 20,000-unit stockpile is almost sufficient if 50% to 60% of the state's stock is usable for COVID-19. In all scenarios, there are shortages, which, in the severe scenario, are on the order of ~28,500 ventilator-days, with a peak shortage of ~2,700 ventilators on 12 April 2020. At its peak, New York would require an additional 7,600 - 9,200 ventilators, and between 400 - 2,000 of those could be donated to other states when New York's peak was over. As seen in figure 2, from the results table of the paper, inter-state sharing reduces the total shortage by 13% to 19% in all the cases.



Data: Mehrotra et al. (2020), Table 1, $\gamma = 75\%$, 23 March to 31 May 2020. Chart drawn by the authors of this report.

Figure 2. Total ventilator shortage with and without sharing (Paper B, $\gamma = 75\%$)

Table 3. Total shortage in ventilator-days by reserve level γ and sharing level τ (Paper B, Table 1)

(γ , τ)	Case I	Case II	Case III	Case IV
(50%, 0%)	112	0	23	303
(50%, 50%)	0	0	0	0
(60%, 0%)	611	96	2,525	2,001
(60%, 50%)	0	0	157	32
(75%, 0%)	4,395	2,877	15,749	28,530
(75%, 50%)	3,668	2,337	13,300	24,828

Values rounded to the nearest ventilator-day. Cases I to IV put increasing probability weight on the top quarter of the IHME forecast range.

The impact is real: a model was provided that allowed planners to conduct rapid allocation policy tests, and an online planning tool was made available by the authors to planners. The paper was awarded the Harold W. Kuhn award for the best paper in Naval Research Logistics.

- Strengths: captures timing (one-off allocation does not); treats uncertainty explicitly, using scenarios; shows value of cooperation between states in numbers; built from public data and runs on standard software.
- There are also several restrictions—results rely on forecasts as well as on approximate estimates of the stocks of goods held in the states; transport times are not taken into account; equipment compatibility is not considered; and the parameter ρ for the safety-stock is fixed; and—most importantly—equality among the states is not an explicit goal and the cost of total shortage may be high enough that one state ends up in much worse shape.
- Extensions: a multistage version that re-plans as data come in (which the authors were working on); equity or worst-state objectives; connection of ventilators to ICU beds and staff (a ventilator without a bed, or a ventilator without staff); and application of to oxygen sharing or ICU-bed sharing between Indian states.

5. Mathematical Model: The Hospital Resource Allocation LP

The two papers share the basic elements: limited resources are used at fixed rates by patient types; capacity must be adhered to. The general model below is stated and then the goal programming form in Paper A is illustrated.

Table 4. Notation for the general model

Symbol	Meaning
$j = 1 \dots n$	Patient groups (departments or case types)
$r = 1 \dots m$	Resources (ward bed-days, ICU bed-days, OT hours, nursing hours, doctor hours)
c_j	Value of one case of type j (contribution margin, patients served or health benefit)
a_{rj}	Amount of resource r used by one case of type j
b_r	Amount of resource r available in the planning period
L_j, U_j	Minimum service level and maximum demand for patient group j
x_j	Decision variable: number of cases of type j planned
d^-, d^+	Under- and over-achievement deviation variables for a goal

Formulation (LP form)

$$\text{Maximise } Z = \sum_{j=1}^n c_j x_j \quad (1)$$

subject to:

$$\sum_{j=1}^n a_{rj} x_j \leq b_r, \text{ for each resource } r = 1 \dots m \quad (2)$$

$$x_j \geq L_j, \text{ for each } j \text{ (minimum service)} \quad (3)$$

$$x_j \leq U_j, \text{ for each } j \text{ (demand ceiling)} \quad (4)$$

$$x_j \geq 0 \text{ (integer if whole patients are required)} \quad (5)$$

Equation (1) maximises the total value of the case mix. Constraint (2) keeps the use of each resource within capacity. Constraint (3) protects essential services such as maternity and emergency medicine, and (4) stops the model planning more cases than will arrive. The dual of this LP attaches a price $y_r \geq 0$ to each resource. At the optimum, y_r is the gain in Z from one more unit of resource r , and a resource with spare capacity has $y_r = 0$. The reduced cost of patient group j is $c_j - \sum_r a_{rj} y_r$; when it is negative, each extra case of that type displaces more valuable work, so the model holds that group at its minimum.

Goal programming form. When there are several targets, each becomes a goal equation with deviation variables:

$$\sum_j c_j x_j + d_0^- - d_0^+ = F \text{ (financial break-even target)} \quad (6)$$

$$x_j + d_j^- - d_j^+ = T_j, \text{ for each } j \text{ (volume targets)} \quad (7)$$

$$\sum_{j \in J(k)} p_j x_j + e_k^- - e_k^+ = I_k, \text{ for each physician } k \text{ (income targets)} \quad (8)$$

$$\text{Minimise } P_1(d_0^-) \gg P_2(\sum_j w_j d_j^-) \gg P_3(\sum_k e_k^- + e_k^+) \quad (9)$$

Constraints (2) to (5) remain as hard limits. In (9), the symbol $\gg$ means the goals are solved in order of priority (pre-emptive goal programming); alternatively the deviations can be added with weights into a single objective. Equation (8), with p_j the physician fee per case and $J(k)$ the cases of physician k , is the kind of goal that Paper A adds to protect physician income. Paper B keeps the same resource-balance logic but adds a time index, inventory balance equations and scenarios for demand.

6. Extension: A Numerical Illustration for a 220-Bed Hospital

In order to demonstrate how the model can be used in practice, we constructed and solved an example case. All the data are assumed, but realistic for a private multi-speciality hospital in an Indian metro city. The hospital has 220 ward beds (85% occupancy), 20 ICU beds (90% occupancy), six operating theatres, operating 10 hours a day, six days a week,

5,400 hours of direct nursing time and 2,200 consultant hours a week. It's a competition between 5 departments for these resources. The data are presented in Table 5.

Table 5. Input data per case and weekly capacity

Department	Contrib. (₹'000)	Ward bed-days	ICU bed-days	OT hours	Nursing hours	Consultant hours	Min / Max cases
General Medicine	12	4.5	0.4	0	18	6	80 / 120
General Surgery	28	5.0	0.6	2.5	22	8	30 / 70
Orthopaedics	35	6.0	0.3	3.0	24	9	15 / 45
Cardiology	60	4.0	1.5	3.5	30	12	10 / 40
Obstetrics & Gynaecology	15	3.0	0.2	1.2	14	5	60 / 80
Weekly capacity		1,309	126	360	5,400	2,200	


There are 5 variables, and 15 constraints for the LP (1-5). It was solved with the HiGHS solver in Python (using SciPy's linprog), and solved once more with integer variables to see the impact of rounding.

6A. Optimal case mix

Table 6. Optimal weekly case mix

Department	LP solution (cases)	Integer solution (cases)	Contribution, (₹'000)	LP
General Medicine	119.40	117	1,432.8	
General Surgery	30.00	30	840.0	
Orthopaedics	43.67	43	1,528.4	
Cardiology	23.43	24	1,405.6	
Obstetrics & Gynaecology	60.00	60	900.0	
Total	276.5	274	6,106.8	

The highest plan gets a weekly payment of ₹61.1 lakh (₹6,106.8 thousand). For planning such a scale, the LP answer is good enough, as the integer plan earns ₹60.9 lakh (which is 0.3% less than the LP answer). Figure 3 displays the location of each department in relation to its designated range. The number of cases in Cardiology and Orthopaedics is increased, while General Medicine is almost at the upper bound of its demand, as there is no theatre time required in this department, and general surgery and obstetrics are at a minimum service level.



Hypothetical data for a 220-bed hospital. Solved with the HiGHS solver in Python (SciPy).


Figure 3. Optimal weekly admissions within each department's allowed range

6B. Bottlenecks and shadow prices

The resources limiting the plan are shown in the figure and table. ICU [intensive care unit] bed-days, OT hours and nursing hours are all fully used. Ward beds (93.4% used) and consultant time (87.8% used) are not constrained so have a shadow price of zero -- meaning that additional units of these would not provide any benefit.

Table 7. Resource use and shadow prices at the optimum

Resource	Used	Available	Utilisation	Shadow price (₹)
Ward bed-days	1,223.0	1,309	93.4%	0
ICU bed-days	126.0	126	100.0%	16,131
OT hours	360.0	360	100.0%	7,588
Nursing hours	5,400.0	5,400	100.0%	308
Consultant hours	1,930.5	2,200	87.8%	0



Shadow price = rise in weekly contribution from one more unit of the resource (₹ per bed-day or per hour).


Figure 4. Capacity used and shadow price of each resource

The shadow prices transform the solution to management counsel. An additional ICU bed-day translates to about ₹16,131 per week and an additional staffed bed (6.3 bed-days per week at 90% occupancy) translates to about ₹1.02 lakh per week, which can then be directly matched with the cost of equipping and staffing the bed. An additional hour

of theatre equates to ₹7,588, allowing for the possibility of theatre hours continuing into the evening. In fact, one additional hour of nursing adds up to just an extra ₹308. The decreased costs is why surgery is kept to a minimum: a General Surgery case has a cost of ₹28,000 but consumes 0.6 of an ICU bed-day, 2.5 theatre hours and 22 nursing hours, which have a combined shadow cost of approximately ₹35,400, meaning that every additional surgery case would reduce contribution by approximately ₹7,400. It does not say anything negative about the department, it simply says that at present capacity, theatre time is better spent in Orthopaedics and Cardiology.

6C. Sensitivity to ICU capacity

Given that the shadow price for the ICU is the highest, we re-solved the model for each size of the ICU between 8 beds and 40 beds, and fixed all the other parameters in the model (Figure 5). If there are less than 13 beds, the problem is infeasible: There are not enough intensive care, even for the minimum service level. The shadow price of the ICU falls between 13 and 23 beds; it is the contribution of a new bed, albeit at a declining rate, as the ICU expands, so does the other resource (theatre time, then nursing), become the binding resource. After 23 beds the curve levels off: additional beds in the ICU would not benefit unless there are additional beds in theatre and nursing beds. A similar test on nursing reveals that the shadow price of ₹308 is valid for up to 5,410 hours per week of nursing, beyond which, nursing becomes non-bottleneck.



Each point is a separate LP run with all other capacities held at base values.

Figure 5. Maximum weekly contribution as the number of staffed ICU beds changes

6D. A funding cut: LP compared with goal programming

The context for Paper A was a funding reduction. In a similar manner, we reduced the nursing hours by 10% (from 5400 to 4860 per week) and introduced two types of goal: Break even contribution of ₹58.0 lakh per week and last year's volume in departments (115, 45, 40, 25, 70 cases). Three plans were compared (Table 8 and Figure 6).

Table 8. Case mix after a 10% cut in nursing hours

Department	Target	LP: contribution max	GP: finance goal first	GP: volume goals first
General Medicine	115	88.2	89.0	85.1
General Surgery	45	30.0	30.0	45.0
Orthopaedics	40	31.0	40.0	25.3
Cardiology	25	34.3	26.4	25.0
Obstetrics & Gynaecology	70	60.0	60.4	70.0
Weekly contribution (₹ lakh)	58.0	59.40	58.00	57.18
Total volume shortfall (cases)	0	60.8	50.7	44.6




Figure 6. Admissions by department under the three plans, against last year's volumes

The plain LP continues to post a positive gain of ₹59.40 lakh, albeit 2.61% less than prior to the cut, but with a 26.1% reduction in General Medicine and a 12.3% increase in Cardiology. The finance-first goal program includes exactly the break-even case, and then some to get Orthopaedics and Cardiology back on track, to the cost of leaving the other two general departments, General Medicine 26.0 cases and Surgery 15.0 cases short. The volume-first goal program returns to last year's level, but falls short of break-even by ₹82,000 per week. There is no plan which achieves all goals: the cut makes a choice, and goal programming shows the price of the choice. The main idea of Paper A is reproduced on a small scale.

7. Reflection

Table 9. Comparison of the approaches

	Paper A (Blake & Carter)	Paper B (Mehrotra et al.)	Our illustration
Decides	Case mix and volume per physician for a hospital division	Daily allocation and return of ventilators across states	Weekly admissions per department
Objective	Minimise deviations from finance, income and practice goals	Minimise expected total shortage	Maximise contribution; then goal programming
Method	Linear goal programming (two linked LPs)	Scenario-based stochastic MILP, Gurobi	LP and MILP with HiGHS; parametric analysis
Uncertainty	Not modelled	Demand scenarios from forecasts	Tested through sensitivity analysis only
Main weakness	Linear costs, subjective weights	No re-planning; equity not modelled	Hypothetical data; single period

7A. Strengths of the modelling approach

The close-to-linear relationships in the LP family are ideal for hospital allocation, and the actual relationships are indeed close-to-linear over the range that matters: twice as many hip replacements require about twice as much theatre time. The models are open, allowing doctors and administrators to review the program inputs; they can be solved in seconds even for large national scales; and the dual values provide price information on beds, theatre time, and staff that the hospital would not otherwise have. With the introduction of goal programming (Paper A), several stakeholders can be considered, and with the introduction of scenario-based stochastic programming (Paper B), a disciplined approach to planning in the face of uncertain demands is added.

7B. Limitations and assumptions

All three models assume that the amount of resources consumed in a patient's stay are fixed and known, and that the length of stay varies greatly from patient to patient, while the demand for resources in the ICU comes in surges. There is a deterministic LP, and on poor days, the hospital that runs at 100% of their planned capacity will be unable to accept patients; queueing theory demonstrates that at 99.9% of their planned capacity, the hospital will have long waiting times. The potential goal also impacts on the service mix – in our example maximum contribution – which pushes out low margin but necessary care – and nothing more than the minimum service constraints can keep it in. Goal programming weights are subjective, and the particular region may not be treated fairly by shortage minimizing models like Paper B. Last but not least, our illustration is based on hypothetical data and does not represent the situation of any real hospital.

7C. Possibilities for extension

Natural extensions include: stochastic or robust case mix models which set a buffer of ICU beds; multi-period models which follow patients from theatre to ICU to ward; integer decisions (open a new ward, open a new theatre); equity or waiting-time objectives in addition to financial ones; and real-time re-planning as suggested by the authors of Paper B. The same structure can be adopted to allocate ICU beds, oxygen and ventilators across district hospitals during peak hours and beds allocated under the Ayushman Bharat PM-JAY scheme between the packages. Our illustration would become evidence if it were validated with real hospital data.

8. Conclusion

How a few lines of linear algebra can help with a decision that matters, in the case of hospitals and their resources. Paper A demonstrates that the problem of a 18% reduction in funding for a hospital can be solved by using linear goal programming to select an optimum case mix that ensures a balanced budget and maintains physicians' incomes and practice. Paper B demonstrates how a stochastic linear model can be used to allocate and share ventilators across a country during a pandemic; and how cooperation between states cut ventilator shortages by 13% to 19%. Two broad principles (the objective over types of patient and the limitations on the capacity of each resource and the limits on service that could be provided) were solved for a hypothetical 220-bed hospital. It highlighted the binding resources as ICU beds, theatre time and nursing; estimated the price of an additional ICU bed at approximately ₹1 lakh per week; proved that the expansion of ICU is not cost effective beyond 23 beds and demonstrated with goal programming that a 10% reduction in nursing is not possible without compromising on financial and volume targets. The key takeaway messages are: (1) shadow prices are frequently as important as the plan itself; (2) hospitals require multi-goal models, not single-objective models; and (3) the most important next step is to model uncertainty in the demand and length of stay explicitly.

References

- Blake, J. T., & Carter, M. W. (2002). A goal programming approach to strategic resource allocation in acute care hospitals. *European Journal of Operational Research*, 140(3), 541-561. [https://doi.org/10.1016/S0377-2217\(01\)00219-3](https://doi.org/10.1016/S0377-2217(01)00219-3)
- Blake, J. T., & Donald, J. (2002). Mount Sinai Hospital uses integer programming to allocate operating room time. *Interfaces*, 32(2), 63-73. <https://doi.org/10.1287/inte.32.2.63.57>
- Charnes, A., & Cooper, W. W. (1961). *Management models and industrial applications of linear programming*. New York: Wiley.
- Dantzig, G. B. (1963). *Linear programming and extensions*. Princeton, NJ: Princeton University Press.
- Hof, S., Fügner, A., Schoenfelder, J., & Brunner, J. O. (2017). Case mix planning in hospitals: A review and future agenda. *Health Care Management Science*, 20(2), 207-220. <https://doi.org/10.1007/s10729-015-9342-2>
- Huangfu, Q., & Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1), 119-142.
- Hulshof, P. J. H., Kortbeek, N., Boucherie, R. J., Hans, E. W., & Bakker, P. J. M. (2012). Taxonomic classification of planning decisions in health care: A structured review of the state of the art in OR/MS. *Health Systems*, 1(2), 129-175.
- Mehrotra, S., Rahimian, H., Barah, M., Luo, F., & Schantz, K. (2020). A model of supply-chain decisions for resource sharing with an application to ventilator allocation to combat COVID-19. *Naval Research Logistics*, 67(5), 303-320. <https://doi.org/10.1002/nav.21905>
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17, 261-272.
- World Bank. (2025). *World Development Indicators*: Hospital beds (per 1,000 people), SH.MED.BEDS.ZS; Physicians (per 1,000 people), SH.MED.PHYS.ZS. Retrieved October 2026 from <https://data.worldbank.org>