Minimax Theory and its Applications

Volume 11 (2026), No. 3s, 151 to 166

Constrained Optimization of Workforce Allocation under Time-Varying Operational Uncertainty

Darshith S.A[1], Rakesh D[2], Aruna Adarsh[3], Tharaka Rami Reddy[4], Aloka K R[5], Vidya D.V[6]

[1]Research scholar , JSS Center for Management Studies , JSS Science and Technology University ,Mysuru Email: darshithsa@jssstuniv.in, Orcid Id : 0009-0009-1072-2078

[2]Associate Professor, JSS-Center for Management Studies, JSS Science & Technology University, JSS TI Campus, SJCE, Mysuru-570006, Email: rakeshdamodar@jssstuniv.in, ORCID: 0000-0003-3963-3174

[3]Professor, JSS-Center for Management Studies, JSS Science & Technology University, JSS TI Campus, SJCE, Mysuru-570006, Email: aruna@jssstuniv.in, ORCID: 0000-0002-4941-8438

[4]Professor, The Oxford College of Engineering, Bengaluru-570006, Email: ktrreddy@gmail.com, ORCID: 0000-0003-4835-4433

[5]Assistant Professor, JSS-Center for Management Studies, JSS Science & Technology University, JSS TI Campus, SJCE, Mysuru-570006, Email: alokakr@jssstuniv.in, ORCID: 0009-0007-7930-9033

[6]Assistant Professor, JSS-Center for Management Studies, JSS Science and Technology University ,Mysuru Email: vidyadv@jssstuniv.in , Orcid Id : 0009-0009-3685-7624

Abstract

Workforce planning under the simultaneous uncertainty and stochasticity of demand requires a balancing act between reliability, labour costs, and nimbleness of operations. Conventional robust workforce planning problems model the uncertainty radius as a parameter, specified externally to the model; however, operational uncertainty propagates as forecast errors unfold. A parsimonious adaptive robust workforce planning model is constructed in which the uncertainty radius, when applied to the current planning horizon, is updated based on the absolute forecast error realized during the immediately preceding period. The proposed mechanism is defined by $\delta_{t+1} = \delta_0 + \gamma \mid D_t - \mu_t \mid$, where $\delta_0$is the baseline uncertainty level and γ controls feedback responsiveness. The resulting workforce problem aims to minimize the costs of the typical workforce, the penalties for shortages, and the costs of transitioning across periods, within workforce capacity constraints. A further contribution to the structure is an exact explicit representation of the solution to the continuous one-period problem; since the cost function is convex and piecewise linear, the optimal workforce level is obtained by evaluating a finite set of candidates: the feasible boundary points, the previous period's workforce, and the demand breakpoint of the robust demand. A contraction result is also established, indicating that the feedback workforce allocation map is contractive.$0 \leq \gamma < 1$ under stationary conditions. Numerical experiments over 120 periods and 500 replications compare low-uncertainty, time-varying-uncertainty, and shock environments. Results show that feedback has little value under stable low uncertainty, but substantially reduces shortage cost and service violations under dynamic and shock-prone uncertainty. In the time-varying environment, $\gamma = 0.60$reduces mean total cost by 3.31%, shortage cost by 58.41%, and service violations by 55.05% relative to the nonadaptive benchmark. In the shock environment, the strongest tested feedback level $\gamma = 0.95$, reduces total cost by 9.73%, shortage cost by 34.44%, and service violations by 46.04%. The results demonstrate that forecast-error feedback provides a mathematically simple and operationally interpretable extension of robust workforce allocation.

Keywords: Workforce Allocation; Robust Optimization; Adaptive Uncertainty; Forecast Error; Time-Varying Demand; Convex Optimization; Staffing; Operational Uncertainty

1. Introduction

Staffing decisions are a fundamental optimization issue across most service operations, including healthcare, call centres, logistics, retail, maintenance, and other labour-dependent systems. The essence of the management problem seems trivially simple: have sufficient workers available to operate at demand, but not so many that excess labour costs accrue. It is made complex when demand exhibits temporal variation and when forecast precision is limited. Hence, traditional workforce planning combines capacity planning with uncertainty management, scheduling constraints, and service-level specifications (Abernathy *et al*., 1973; Ernst *et al*., 2004).

The body of the literature comprises a large set of approaches to workforce planning. Early mathematical formulations viewed workforce planning asa multistage decisionproblem,where later work decisions depend on information that becomes available only after earlier decisions are made (Dantzig 1955). Recent developments in service-sector manpower planning have combined stochastic programming with chance constraints (Abernathy *et al*. 1973). In the last decade, researchers have studied dynamic staffing, multiskilling, routing, stochastic demand, and robust optimization. As in survey articles, workforce scheduling has become a well-established operations research area with numerous formulations for different application scenarios (Ernst *et al*. 2004; Castillo-Salazar *et al*. 2012).

The aspect of uncertainty receives a particularly significant part. Demand in service systems is typically nonstationary; variations in its magnitude and in its variability can appear over time. Nonstationary demand creates, as described by Defraeye and Van Nieuwenhuyse (2016), a direct connection between forecasting, staffing and the performance of the system. Such issues are similar for call centres, where staffing must be adapted whenever arrival rates and demands change (Gans *et al*., 2003; Cezik & L'Ecuyer, 2008).

For this kind of problem, however, robust optimization represents an attractive methodological approach since there is no need for a fully specified probability distribution for the uncertain data (Ben-Tal & Nemirovski, 1998; Ben-Tal *et al*., 2009). Uncertainty is instead described by means of uncertainty sets, and the decisions are taken so as to perform well for all realizations within this set (Ben-Tal & Nemirovski, 1998). The trade-off between the nominal solution and a protected one is called the price of robustness (Bertsimas & Sim, 2004). Hence, robust optimization has been recently used as an important device for planning under partial information (Bertsimas *et al*., 2011).

However, there is one practical limitation. A number of workforce planning models have shown that the value of the uncertainty radius can be determined at the outset of a planner's operations and remains fixed over the planning horizon. In practice, the level of uncertainty may be read from the magnitude of errors in a forecast.

When the demand is much greater than the forecast for a particular period, then the next period can be modelled as an uncertain one; however, when the error is not large, then there would be no need to carry an excessive amount of worker overcapacity. The current paper develops a simple feedback mechanism that can address this issue. Rather than assuming a constant radius of uncertainty, the model proposed here updates the degree of uncertainty based on the absolute forecast error of the last period.

The model is, in essence, also extremely parsimonious. This paper does not seek to propose a new general theory of endogenous uncertainty, since decision-dependent, adaptive uncertainty has long been considered a well-known notion within robust optimization (Bertsimas *et al*., 2011). Instead, this paper provides a specific mechanism for the evolution of forecast-error-driven uncertainty in workforce allocation problems, as well as an exact solution expression, a stability result and numerical evidence of the responsiveness-cost trade-off.

2. Literature Review and Research Gap

2.1 Workforce scheduling and allocation

Workforce scheduling has long been a topic with roots in the history of operations research. Dantzig (1955) provides an origin for the viewpoint that decisions are made under uncertainty, since subsequent decisions require information to be revealed only afterwards. Abernathy *et al*. (1973) extended this reasoning for the context of multistage workforce planning and scheduling in service organizations. The formulation that they proposed emphasizes the fundamental conflicts that occur between the workforce decisions that have been made early on and requirements that are uncertain in the future.

The subsequent literature expanded significantly. Ernst *et al*. (2004) reviewed models of staff scheduling and rostering, noting the importance of service demands, flexible working schedules, preferences, and operational constraints in such models. Baker (1976) presented a review of workforce allocations in response to cyclic demand, showing how mathematical models can harness cyclic demand patterns. Castillo-Salazar *et al*. (2012) highlighted the scope of workforce scheduling problems, particularly when task constraints are coupled with workforce assignments and geographical routing.

Context for application: Service systems provide an important application environment. Gans *et al*. (2003) defined the tight coupling between demand, staff, queueing performance and service quality at a call centre. Cezik and L'Ecuyer (2008) illustrated how to combine staffing decisions with simulation-based estimates of

service levels for a call centre with multiple skills. These studies all highlight the criticality of making workforce decisions that are sensitive to the stochastic operational environment.

2.2 Optimization under uncertainty

In the framework of stochastic programming, data uncertainty is explicitly modelled through probability distributions and scenario structures. According to Birge and Louveaux (2011), stochastic programming is an optimization framework in which decisions are made in the presence of data uncertainty, often, but not necessarily, through multistage or two-stage models. Such models can be very effective if reliable probability distributions can be derived.

The field of robust optimization provides a perspective that complements other lines of research. While Ben-Tal and Nemirovski (1998) laid down the mathematical foundation for robust convex optimization, Ben-Tal *et al*. (2009) gave a comprehensive and systematic treatment of the subject for the case of data uncertainty. Bertsimas and Sim (2004) highlighted the importance of controlling the cost of robustness rather than simply seeking the maximum possible level of protection. More recently, Bertsimas *et al*. (2011) connected robust optimization to adaptive and multistage decision making.

The appeal of robust optimization is high when, at least for the allocation of workforce resources, uncertainty can be accounted for without the need for a full specification of the demand distribution. However, when forecast quality fluctuates over time, a fixed uncertainty radius may become inappropriate. A radius that provides enough cushion during a stable period might not be enough following a sudden shock.

2.3 Recent workforce uncertainty models

According to current research, this area is growing towards a flexible, data-driven, and self-adaptive workforce optimization. Wang *et al.* (2025) presented a distributionally robust flexible personnel scheduling based on Wasserstein ambiguity sets, which addresses demand uncertainty and the lack of knowledge about the actual distribution. The approach that Wang *et al.* Proposed focused on the performance under the worst-case expectation along with a flexible personnel schedule.

Rath *et al.* (2025) have developed a multistage robust structure for staff planning in a healthcare system, where patient assignment, on-call deployment and operational deployment are linked by demand uncertainty. Their paper demonstrates the advantages of dynamic staffing and well-calibrated uncertainty sets. Many other recent studies have considered robust workforce scheduling against stochastic demand. Hu *et al.* (2024) studied multi-objective robust emergency staff scheduling, and Porto *et al.* (2025) proposed workforce flexibility through annualized hours, multi-skilling and overtime, alongside stochastic demand.

The results from these studies show that both the robust and stochastic methodologies for workforce optimization are well cited in the literature; thus, the incremental development of a robust workforce methodology with time-varying demand is not a sufficiently novel contribution.

2.4 Research gap

The specifically chosen lacuna, however, of this study is the evolution of the uncertainty radius. Standard robust workforce approaches derive uncertainty from the past, empirical distributions, or some calibration routine. By comparison, the proposed approach relies on a unique, recursive operational principle:

$$\delta_{t+1} = \delta_0 + \gamma \mid D_t - \mu_t \mid.$$

There is an intuitive feel to this logic: a significant forecasting error thus becomes information about future operational uncertainty. The model then does not model uncertainty as static, and does not require a future complex distributionally robust specification. This rather conservative approach would not be sufficient or feasible for the dissertation: the present paper does not make any claim about using endogenous uncertainty, robust workforce planning, or adaptive optimization as new knowledge.

3. Adaptive Robust Workforce Allocation Model

3.1 Notation

Consider planning periods $t = 1, \ldots, T$. Let:

- x_t: workforce allocated in period t;
- W_t: maximum available workforce;
- D_t: realized demand;
- μ_t: forecast demand;
- $a > 0$: productivity per worker;
- δ_t: adaptive uncertainty radius;
- U_t: robust upper demand;
- $c > 0$: regular workforce cost per worker;
- $p > 0$: shortage penalty per unit of unmet demand;
- $r \geq 0$: workforce adjustment cost;
- $\gamma \geq 0$: forecast-error feedback coefficient.

The workforce capacity constraint is

$$0 \leq x_t \leq W_t.$$

The effective capacity supplied by the workforce is ax_t. A shortage occurs whenever $D_t > ax_t$.

3.2 Forecast-error feedback

Define the forecast error in period.$t$as

$$e_t = D_t - \mu_t.$$

The adaptive uncertainty radius is

$$\boxed{\delta_{t+1} = \delta_0 + \gamma \mid e_t \mid.} \qquad (1)$$

The absolute value is used because both positive and negative forecast errors provide information about uncertainty magnitude. A forecast error of $-15$is is interpreted as evidence of uncertainty of the same magnitude as a forecast error of $+15$.

The robust upper demand for the next period is

$$\boxed{U_{t+1} = \mu_{t+1} + \delta_0 + \gamma \mid D_t - \mu_t \mid.} \qquad (2)$$

When $\gamma = 0$,

$$U_t = \mu_t + \delta_0,$$

which recovers the fixed-radius benchmark.

Thus, the proposed model is nested:

$$\boxed{M_{\gamma=0} = M_{\text{base}}.} \qquad (3)$$

This nesting is important because it permits a direct comparison between adaptive and nonadaptive uncertainty.

3.3 Objective function

The total operational cost is

$$\boxed{\min_{x_1,\ldots,x_T} \sum_{t=1}^{T} [cx_t + p(U_t - ax_t)^+ + r \mid x_t - x_{t-1} \mid],} \qquad (4)$$

where

$$z^+ = \max(z, 0)$$

The first element captures the costs of labour. The second element penalizes the firm when capacity is less than the high demand. The third element captures the cost of flexibility of the workforce.

The identified model incorporates three effects that work in opposition.When x_t increases, costs on the workforce are higher. However, the more workers, the less the exposure to shortages; likewise, high mobility lowers the risk of shortages, but increases the cost of mobility. The formulation is purposely kept simple so as to preserve the logical framework of the problem while still capturing the main feature of the operational environment: uncertainty.

4. Mathematical Analysis

4.1 Convexity of the one-period problem

Given $x_{t-1} = v$, the one-period optimization problem is

$$\min_{0 \le x \le W} f(x) = cx + p(U - ax)^+ + r \mid x - v \mid. \qquad (5)$$

The function $cx$is linear. The shortage function $p(U - ax)^+$is convex because it is the maximum of two affine functions, and $r \mid x - v \mid$is convex. Therefore $f(x)$is convex.

Proposition 1. The one-period adaptive workforce problem is a convex piecewise-linear optimization problem.

Proof. Each component of (5) is convex. The positive-part function is the maximum of $U - ax$and zero, and therefore is convex. The absolute-value function is convex. The sum of convex functions is convex. Because all terms are piecewise linear, $f(x)$is piecewise linear and convex. This property is significant because no general-purpose nonlinear optimizer is required.

4.2 Exact candidate-set solution

Define the robust demand breakpoint.

$$q = \frac{U}{a}.$$

The only points at which the slope of the objective can change are:

1. $x = 0$;
2. $x = W$;
3. $x = v = x_{t-1}$;
4. $x = q = U/a$.

Consequently, define the candidate set.

$$\mathcal{C} = \{0, W, \text{clip},(v, 0, W)\text{clip}(q, 0, W)\}. \qquad (6)$$

Proposition 2. An optimal solution of (5) is attained at an element of $\mathcal{C}$.

Proof. Between any two consecutive points in $\mathcal{C}$, both the positive-part term and the absolute-value term have fixed linear forms. Therefore $f(x)$is affine on each interval. An affine function on a closed interval attains its minimum at an endpoint. The endpoints of all such intervals belong to $\mathcal{C}$. Hence, at least one global minimizer belongs to $\mathcal{C}$. □

The exact solution is therefore

$$\boxed{x_t^* \in \arg\min_{x \in \mathcal{C}_t} [cx + p(U_t - ax)^+ + r \mid x - x_{t-1} \mid].} \qquad (7)$$

This result is one of the main mathematical advantages of the proposed model. The adaptive uncertainty mechanism does not destroy tractability.

For integer workforce decisions, the same principle can be used by evaluating the integer neighbours of the breakpoints. Thus, the candidate set can be expanded to include $\lfloor q \rfloor, \lceil q \rceil, \lfloor v \rfloor, \lceil v \rceil$, together with 0 and W.

4.3 Special case without adjustment cost

If $r = 0$, the problem reduces to

$$\min_{0\le x\le W} cx + p(U - ax)^+.$$

For $x < U/a$, the slope is $c - pa$. For $x > U/a$, the slope is c.

Therefore, when

$$pa > c, \qquad (8)$$

the objective decreases before $U/a$and increases afterwards. Hence

$$\boxed{x_t^* = \min\left\{W_t, \frac{U_t}{a}\right\}.} \qquad (9)$$

Substituting the adaptive uncertainty rule gives

$$\boxed{x_t^* = \min\left\{W_t, \frac{\mu_t + \delta_0 + \gamma|D_{t-1} - \mu_{t-1}|}{a}\right\}.} \qquad (10)$$

Equation (10) clearly decomposes workforce requirements into three components:

$$x_t^* = \frac{\mu_t}{a} + \frac{\delta_0}{a} + \frac{\gamma \mid e_{t-1} \mid}{a},$$

provided the capacity bound is inactive.

The first component is forecast-based staffing, the second is the baseline robustness buffer, and the third is an adaptive error-response buffer.

4.4 Stability of the feedback mechanism

Consider a stationary setting with a constant forecast.μ, demand D, productivity a, and workforce capacity W. The simplified adaptive mapping is

$$F(x) = \min\left\{W, \frac{\mu + \delta_0 + \gamma|D - ax|}{a}\right\}. \qquad (11)$$

The absolute-value inequality

$$|| D - ax \mid - \mid D - ay \;|| \le a \mid x - y \mid \qquad (12)$$

implies

$$\mid F(x) - F(y) \mid \le \gamma \mid x - y \mid. \qquad (13)$$

Theorem 1. If $0 \le \gamma < 1$, then $F$is a contraction and has a unique fixed point.

Proof. Equation (13) shows that the Lipschitz constant of $F$is at most γ. For $0 \le \gamma < 1$, the mapping is a contraction. By the Banach fixed-point theorem, a unique fixed point exists and successive application of Fconverges to that point.

Consequently,

$$\boxed{\mid x_t - x^* \mid \le \gamma^t \mid x_0 - x^* \mid.} \qquad (14)$$

This result provides a straightforward quantitative interpretation of the feedback coefficient. In the stationary simplified model, magnitudes less than one imply a stable feedback process, while magnitudes close to one imply a slower contraction along with greater responsiveness.

4.5 Sensitivity to forecast error

When the upper capacity constraint is inactive, and the previous forecast error is positive,

$$D_{t-1} > \mu_{t-1},$$

equation (10) becomes

$$x_t^* = \frac{\mu_t + \delta_0 + \gamma(D_{t-1} - \mu_{t-1})}{a}.$$

Therefore,

$$\boxed{\frac{\partial x_t^*}{\partial \gamma} = \frac{D_{t-1} - \mu_{t-1}}{a}.} \qquad (15)$$

A larger positive forecast error therefore produces a larger marginal benefit from increasing γ. This analytical result anticipates the simulation findings: In settings characterized by substantial and enduring forecast inaccuracies, adaptive feedback ought to prove more beneficial than it would in settings where demand remains steady.

5. Computational Experiment

5.1 Experimental design

The numerical experiment evaluates the proposed model over $T = 120$ planning periods with $R = 500$independent replications for each parameter combination. The workforce productivity is

$$a = 10,$$

regular workforce cost is

$$c = 1,$$

shortage penalty is

$$p = 8,$$

adjustment cost is

$$r = 0.5,$$

maximum workforce is

$$W = 16,$$

and baseline uncertainty is

$$\delta_0 = 10.$$

The demand forecast follows a sinusoidal seasonal pattern:

$$\mu_t = 105 + 18\sin\left(\frac{2\pi t}{T}\right). \qquad (16)$$

This produces a systematic time-varying demand environment while retaining a transparent analytical structure.

Three uncertainty environments are considered.

Scenario A: Low uncertainty. Demand noise has standard deviation 2:

$$D_t = \mu_t + 2\epsilon_t.$$

Scenario B: Time-varying uncertainty**.** The standard deviation evolves as

$$\sigma_t = 6\left[0.55 + 0.75\left|\sin\left(\frac{2\pi t}{T} + 0.4\right)\right|\right],$$

and

$$D_t = \mu_t + \sigma_t \epsilon_t.$$

Scenario C: Shock environment. The same time-varying uncertainty is supplemented by positive demand shocks. In each period, a shock occurs with probability 0.07. Conditional on occurrence, the shock magnitude follows a gamma distribution with shape 2 and scale 10.

In all cases,

$$\epsilon_t \sim N(0,1)$$

The feedback coefficient is varied over

$$\gamma \in \{0, 0.1, 0.2, 0.4, 0.6, 0.8, 0.95\}$$

The value $\gamma = 0$ is the nonadaptive benchmark.

5.2 Performance measures

Five principal performance indicators are reported:

1. **Total cost**, consisting of workforce, shortage, and adjustment costs.
2. **Workforce cost**, equal to $c\sum_t x_t$.
3. **Shortage cost**, equal to $p\ \sum_t (D_t - ax_t)^+$.
4. **Adjustment cost**, equal to $r\sum_t |\, x_t - x_{t-1}\,|$.
5. **Service violation rate**, defined as the percentage of periods for which $D_t > ax_t$.

Two additional indicators are reported: average workforce and uncertainty amplification,

$$A_t = \frac{\delta_t}{\delta_0}. \qquad (17)$$

The exact candidate-set solution of Proposition 2 is used for each period. As a result, while a simple staffing rule is used, the adjustment costs that are reported are for decisions that explicitly take into account the inter-period adjustment term in the optimization.

6. Numerical Results

6.1 Overall cost behaviour

Table 1. Mean performance across 500 replications and 120 periods

Scenario	γ	Total cost	Workforce cost	Shortage cost	Adjustment cost	Violation (%)	Avg. workforce	Amplification
Low uncertainty	0.00	1389.30	1380.00	0.00	9.30	0.000	11.500	1.000
Low uncertainty	0.20	1393.34	1383.78	0.00	9.56	0.000	11.531	1.031
Low uncertainty	0.40	1397.99	1387.60	0.00	10.39	0.000	11.563	1.063
Low uncertainty	0.60	1402.94	1391.37	0.00	11.57	0.000	11.595	1.095
Low uncertainty	0.80	1408.07	1395.15	0.00	12.91	0.000	11.626	1.126

Low uncertainty	0.95	1412.17	1398.14	0.00	14.03	0.000	11.651	1.151
Time-varying	0.00	1557.07	1380.00	167.76	9.30	5.462	11.500	1.000
Time-varying	0.20	1528.47	1391.73	124.93	11.82	4.253	11.598	1.098
Time-varying	0.40	1515.29	1403.47	95.58	16.24	3.383	11.696	1.196
Time-varying	0.60	1505.57	1415.08	69.77	20.72	2.455	11.792	1.292
Time-varying	0.80	1515.12	1426.89	62.59	25.64	2.138	11.891	1.391
Time-varying	0.95	1517.98	1435.62	53.15	29.22	1.843	11.964	1.464
Shock	0.00	2321.97	1380.00	932.67	9.30	10.372	11.500	1.000
Shock	0.20	2254.70	1394.32	846.67	13.71	8.625	11.619	1.119
Shock	0.40	2170.23	1408.65	741.59	19.99	7.307	11.739	1.239
Shock	0.60	2165.74	1422.99	715.97	26.78	6.628	11.858	1.359
Shock	0.80	2136.15	1437.08	665.91	33.16	6.013	11.976	1.479
Shock	0.95	2096.14	1446.92	611.46	37.76	5.597	12.058	1.565

The complete simulation results are reported in Table 1.raises the overall expense, given that the benchmark already circumvents shortages. Consequently, supplementary feedback generates costs related to workforce and adjustment while yielding no equivalent decrease in shortage-related expenditure.

Under time-varying uncertainty, the cost curve is nonmonotonic. Total cost decreases from 1557.07 at $\gamma = 0$to 1505.57 at $\gamma = 0.60$, after which it increases. Thus, moderate feedback provides the best observed cost balance in this experimental setting.

The shock environment exhibits a different pattern. Total cost continues to decline as feedback increases, reaching 2096.14 at $\gamma = 0.95$. This illustrates the core selling point of the model, viz. That a more aggressive feedback response can counterbalance the additional labour force and adjustment costs if demand shocks are sufficiently expensive.

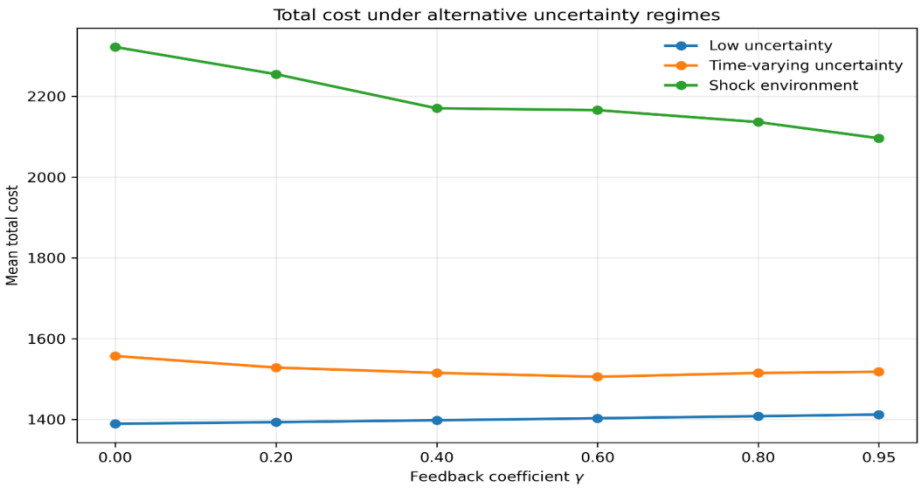


Figure 1. Total cost under alternative uncertainty regimes.

As shown in Figure 1, the economically appropriate feedback level depends on the uncertainty regime rather than being universally large or small.

6.2 Shortage reduction

Figure 2. is concerned with the shortage cost. The shortage cost is negligible with little uncertainty at all levels of feedback, resulting in no benefit from the feedback mechanism in terms of shortages in this setting.

The time-varying environment produces a substantial reduction. Relative to $\gamma = 0$, the shortage cost falls from 167.76 to 69.77 at $\gamma = 0.60$, a reduction of approximately 58.41%. The corresponding total cost decreases by approximately 3.31%.

The shock environment shows a reduction from 932.67 to$\gamma = 0$to 611.46 at $\gamma = 0.95$, representing a 34.44% reduction. The reduction is smaller proportionally than in the time-varying scenario because shock demand can exceed the adaptive upper bound by a large amount. Nevertheless, the absolute shortage-cost reduction is substantial.

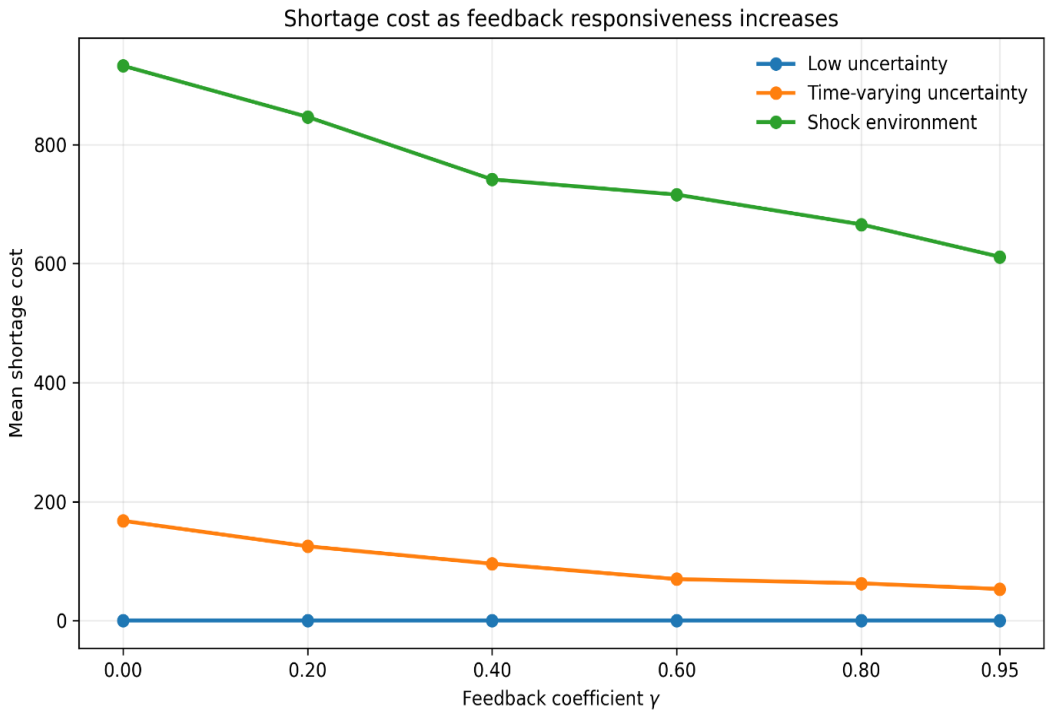


Figure 2. Shortage cost as feedback responsiveness increases.

The inward radial configuration of Figure 2 substantiates the sensitivity result of (15). The (expected) feedback benefit increases as forecast errors increase, since the adaptive uncertainty radius also grows in response to larger forecast errors.

6.3 Service violations

Service violation rates provide an alternative interpretation of the results. As shown in Figure 3, the time-varying environment falls from 5.462% at $\gamma = 0$to 2.455% at $\gamma = 0.60$, representing a reduction of 55.05%.

In the shock environment, the violation rate declines from 10.372% to 5.597% at $\gamma = 0.95$, a reduction of 46.04%.

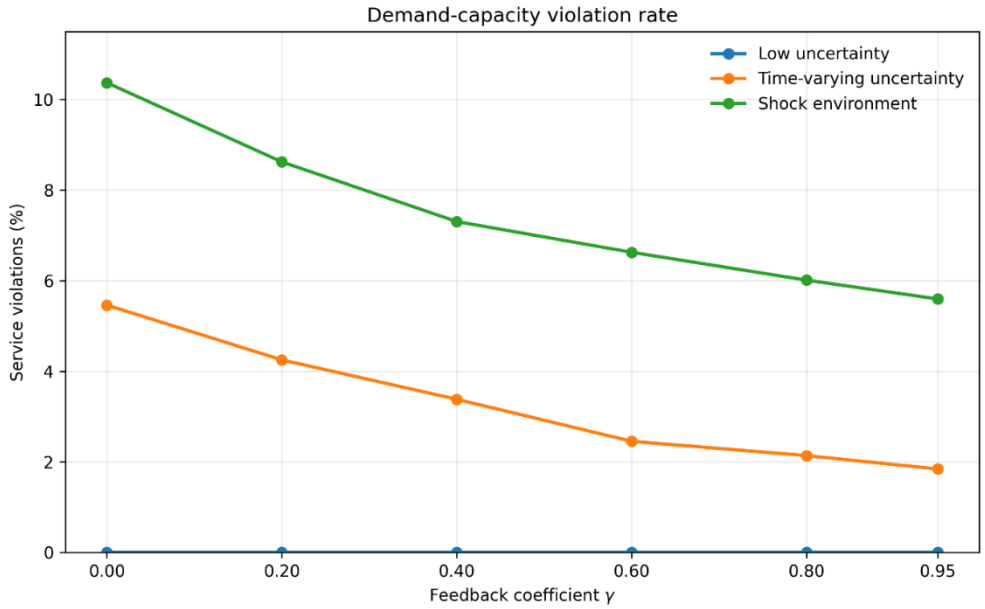


Figure 3. Demand-capacity violation rate.

The results demonstrate that adaptive feedback is particularly relevant when service reliability is operationally important. The reduction in violation frequency occurs because a large error in period $t-1$automatically increases the robust staffing target in period t.

6.4 Workforce and adjustment effects

The shortage reduction is not free. Average workforce increases from 11.50 at $\gamma = 0$to 11.79 at $\gamma = 0.60$in the time-varying environment. At $\gamma = 0.95$, average workforce reaches 11.96.Adjustment cost also increases with feedback. In the time-varying environment, it increases from 9.30 at $\gamma = 0$to 20.72 at $\gamma = 0.60$and 29.22 at $\gamma = 0.95$.The same effect is stronger in the shock environment. Adjustment cost rises to 37.76 at $\gamma = 0.95$. Thus, stronger feedback does not simply increase staffing; it also causes the workforce trajectory to react more strongly to newly observed forecast errors.This is one of the features of the model. If policy responds excessively to each forecasting error, it can generate instability in operations – even if it improves the reduction in shortage exposure. Therefore, the adjustment term is a kind of regularizer that prevents rapid changes in workforce.

6.5 Uncertainty amplification

Figure 4 . presents the average uncertainty amplification factor. By construction,

$$A_t = 1 + \frac{\gamma \mid e_{t-1} \mid}{\delta_0}.$$

Amplification is approximately 1.15 under$\gamma = 0.95$under low uncertainty, 1.46 under time-varying uncertainty, and 1.57 under the shock environment.

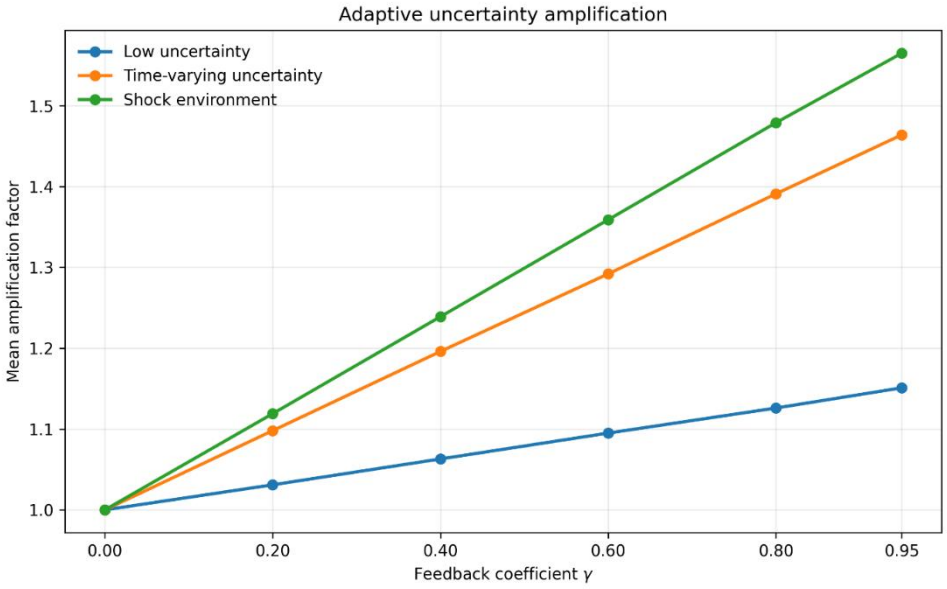


Figure 4. Adaptive uncertainty amplification.

The scenario differences are informative. The same value of γcreates substantially more uncertainty amplification when forecast errors are larger. Therefore, γ should be interpreted as a sensitivity parameter rather than as a direct measure of uncertainty itself.

6.6 Summary of scenario-dependent performance

Table 2 summarizes the benchmark-to-selected-feedback comparisons.

Table 2. Benchmark comparison

Scenario	Benchmark Λ	Selected tested Λ	Total-cost change	Shortage-cost change	Violation-rate change
Low uncertainty	0	0	—	—	—
Time-varying uncertainty	0	0.6	−3.31%	−58.41%	−55.05%
Shock environment	0	0.95	−9.73%	−34.44%	−46.04%

The values selected in Table 2 should not be interpreted as feedback coefficients optimal in general. For the particular experimental design used here, they are the tested values that do best in terms of total cost. The preferred value might change if the costs of shortages, adjustment costs, forecast quality, or the distributions of shocks change.

7. Result & Discussion

The numbers have implications for the problem studied. For example, adaptive uncertainty is most useful in environments where the true uncertain condition changes rapidly. Second, given that there is no uncertainty or limited uncertainty, the nonadaptive model should be able to cope. There is a typical price-of-robustness pattern where the more robust needs are more expensive, even if the quality-of-service increases (γ) little.

Second, the adaptive feedback is valuable when forecast errors carry information about operational risks that could surface later. In the time-varying setting, a modest amount of feedback is expected to bring down shortage costs enough to compensate for the additional costs of workforce and adjustment. This simplistic empirical interpretation of the feedback mechanism is thus that, rather than trying to predict demand, the approach adapts to the prediction performance as it is observed.

Third, the shock environment clarifies why adaptive mechanisms may be desirable in operational settings where demand distributions may be difficult to specify. Distributionally robust methods are always viable alternatives

when historical data exists. The proposed model minimizes data requirements as it requires only the current forecast and the most recently observed error, thus making the mechanism attractive in settings where data is sparse or where the operational environment evolves faster than a statistical model can adapt.

Fourth, the adjustment-cost term is essential. Without $r \mid x_t - x_{t-1} \mid$, the model would react directly to every change in robust demand. In practice, hiring, reassignment, By using more hours of overtime, renting extra temporary workers, and shifting the hours of the rotating workers, the adjustment cost is expected to be incurred. As a result, the adjustment penalty seems to transform the model from a static robust staffing rule into a simple dynamic model of the labor force.

The specific result pertaining to the candidate set is probably of limited practical consequence. The converse is true on a grander scale: in a large-scale scenario, generating the optimal workforce could be computationally infeasible with any number of skills, locations, time windows, and individual employee constraints. The construction in this paper shows that an intentional decision was made not to consider these complexities but instead to attempt to construct a mathematically explicit and intelligible "building block." For the single-period decision, only a relatively small number of candidates need to be considered. Thus the policy might be appropriate for real-time or rolling-horizon implementation.

The contraction result provides another useful interpretation. For $\gamma < 1$, the simplified feedback mapping is stable.This analysis should not be taken to mean that all stochastic implementations are globally stable in the "strongest" sense, since demand, forecasts, capacity constraints, and adjustment costs can all influence this conclusion. Instead, the result simply indicates a strong, sufficient condition for the stationary feedback mapping.

These results also add to the growing body of work connecting adaptive and multistage robust optimization. Robust optimization has been steadily shifting away from static "protection" forms toward decision rules that use information as it becomes available. The model constructed in this paper may be seen as a rather minimalist adaptive decision rule where the state variable is the previous forecast error.

The importance of managing uncertainty dynamically has been recognized in the recent literature on workforce management. There is evidence to suggest that dynamic planning of staffing could have operational benefits if staffing decisions were conditioned on demand forecasts and information obtained later. There is also evidence that labor market flexibility can enhance workforce robustness in the presence of demand uncertainty. By analyzing a mathematically minimal feedback structure, as opposed to a large mixed-integer workforce planning system, this paper complements these lines of research.However, there are a number of important limitations.

Sinusoidal forecasts coupled with some pre-specified noise structures offer high control; however, no external validity is established for any particular industry.In addition, the second point suggests that the feedback rule in absolute-error form does not allow for distinction between systematic bias and stochastic noise. Systematic upward or downward trends in forecast bias may result in the rule providing signals of uncertainty, but not the problem of the forecast bias.However, third, the most important limitation is that there is only one aggregated workforce with the same productivity. Real firms are likely to have multiple types of employees, multiple skills, varied locations, multiple shift structures, and contractual constraints.

Finally, the current form of the model is one step ahead, with uncertainty being updated from the previous period. Many other models are capable of using exponentially weighted historical errors, for instance, to improve performance.

$$\delta_{t+1} = \delta_0 + \gamma \sum_{j=1}^{L} \omega_j \mid e_{t+1-j} \mid,$$

where $\sum_j \omega_j = 1$. Such a formulation would allow the uncertainty radius to depend on a recent history rather than only the most recent error.

Fifth, the present experiments use a fixed shortage penalty. In applications where service-level targets are more important than monetary shortage penalties, a chance-constrained or risk-measure formulation could be considered. Risk measures such as CVaR provide a natural extension for explicitly controlling tail exposure

These limitations also define opportunities for future work. A data-driven estimation method could be developed for selecting γ Instead of looking at a single grid, historical data could be used to define a cost out-of-sample

criterion. This would imply that positive demand surprises should be dealt with differently than negative forecast errors. This modification should be particularly appropriate in situations where the consequences of underforecasting incur greater facility damage than those of overforecasting. A possible extension would be to embed the rule in a distributionally robust framework. The feedback rule could be used to compute the centre or radius of a Wasserstein ambiguity set, resulting in a hybrid model built on error-driven adaptive uncertainty and distributionally robust workforce scheduling.

8. Conclusion

This paper presents an adaptive robust model of workforce planning based on forecast errors, which is intended to handle time-varying operational uncertainty. When the magnitude of the forecast errors is large, the uncertainty buffer for the next period should also be significantly larger, and conversely, when the forecast errors are small, the buffer should be close to the minimum level. Within the context of the workforce planning problem, the model incorporates three costs – labour, shortage, and transition costs – and, although it offers a dynamic mechanism for managing operational uncertainty, the single-period problem remains convex and piecewise linear. The study therefore shows that the optimal solution can be obtained exactly by testing a finite set of candidate workforces, with the added practical benefit over more computationally-demanding robust workforce formulations.The mathematical analysis also establishes a contraction property for the simplified adaptive mapping when $0 \leq \gamma < 1$. This offers a preliminary look at stability with respect to the feedback and helps to explain the logic as to why overly aggressive feedback might produce excessive responsiveness. The simulation experiment shows that there is a major regime-dependent effect. When the level of uncertainty is low, feedback raises costs, while for the stationary level of uncertainty, the baseline policy staves off shortages$\gamma = 0.60$, total cost falls by 3.31%, shortage cost falls by 58.41%, and service violations fall by 55.05% relative to the nonadaptive benchmark. Under the shock environment, stronger feedback produces the largest total-cost improvement among the tested configurations, with $\gamma = 0.95$reducing total cost by 9.73% and service violations by 46.04%.The critical point, then, is not that a certain level of protection is always optimal. Rather, the usefulness of adaptive feedback depends on the underlying mode of operational uncertainty. In a more stable environment, it is possible that protection with very limited feedback is sufficient. But shock-prone and volatile environments should produce substantially higher expected benefits from adaptive protection. This is consistent with the general principle that the cost of robustness should be consistent with the system's real level of uncertainty.

Future studies could experimentally test this mechanism using real- world workforce data, statistically estimate feedback parameters, including asymmetric forecast errors, and generalize the model to multiple groups of workers in various geographical locations. Another potential direction for future research might be to combine this feedback mechanism with distributionally robust optimization. Overall, the methodology in this paper indicates how forecasting can be related to workforce optimization, where forecast error can induce operational uncertainties, then workforce protection, then balance it out with the costs of adaptation.

References

1. Abernathy, W. J., Baloff, N., Hershey, J. C., & Wandel, S. (1973). A three-stage manpower planning and scheduling model—A service-sector example. *Operations Research, 21*(3), 693–711. https://doi.org/10.1287/opre.21.3.693

2. Baker, K. R. (1976). Workforce allocation in cyclical scheduling problems: A survey. *Journal of the Operational Research Society, 27*(1), 155–167. https://doi.org/10.1057/jors.1976.30

3. Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.

4. Ben-Tal, A., & Nemirovski, A. (1998). Robust convex optimization. *Mathematics of Operations Research, 23*(4), 769–805. https://doi.org/10.1287/moor.23.4.769

5. Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review, 53*(3), 464–501. https://doi.org/10.1137/080734510

6. Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations Research, 52*(1), 35–53. https://doi.org/10.1287/opre.1030.0065

7. Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming* (2nd ed.). Springer. https://doi.org/10.1007/978-1-4614-0237-4

8. Castillo-Salazar, J. A., Landa-Silva, D., & Qu, R. (2012). A survey on workforce scheduling and routing problems. In *Proceedings of the 9th International Conference on the Practice and Theory of Automated Timetabling* (pp. 283–302).

9. Cezik, M. T., & L'Ecuyer, P. (2008). Staffing multiskill call centers via linear programming and simulation. *Management Science, 54*(2), 310–323. https://doi.org/10.1287/mnsc.1070.0824

10. Dantzig, G. B. (1955). Linear programming under uncertainty. *Management Science, 1*(3–4), 197–206. https://doi.org/10.1287/mnsc.1.3-4.197

11. Defraeye, M., & Van Nieuwenhuyse, I. (2016). Staffing and scheduling under nonstationary demand for service: A literature review. *Omega, 58*, 4–25. https://doi.org/10.1016/j.omega.2015.04.002

12. Ernst, A. T., Jiang, H., Krishnamoorthy, M., & Sier, D. (2004). Staff scheduling and rostering: A review of applications, methods and models. *European Journal of Operational Research, 153*(1), 3–27. https://doi.org/10.1016/S0377-2217(03)00095-X

13. Gans, N., Koole, G., & Mandelbaum, A. (2003). Telephone call centers: Tutorial, review, and research prospects. *Manufacturing & Service Operations Management, 5*(2), 79–141. https://doi.org/10.1287/msom.5.2.79.16071

14. Hu, Y., Liu, Q., Li, S., & Wu, W. (2024). Multi-objective robust optimization of staff scheduling for emergency under stochastic demand. *Expert Systems with Applications, 254*, 124214. https://doi.org/10.1016/j.eswa.2024.124214

15. Porto, A. F., Lusa, A., Herazo, S. A., & Henao, C. A. (2025). Improving the robustness of retail workforce management with a labor flexibility strategy and consideration of demand uncertainty. *Operations Research Perspectives, 15*, 100345.

16. Rath, S., Rajaram, K., Hudson, M. E., & Mahajan, A. (2025). Multilocation, dynamic staff planning for a healthcare system: Methodology and application. *Operations Research, 74*(1), 199–223. https://doi.org/10.1287/opre.2023.0438

17. Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk, 2*(3), 21–41.

18. Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2014). *Lectures on stochastic programming: Modeling and theory* (2nd ed.). Society for Industrial and Applied Mathematics.

19. Wang, Z., Luo, Z., & Shen, H. (2025). Data-driven robust flexible personnel scheduling. *Computers & Operations Research, 176*, 106935. https://doi.org/10.1016/j.cor.2024.106935

20. Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review, 53*(3), 464–501. https://doi.org/10.1137/080734510

21. Ben-Tal, A., & Nemirovski, A. (1998). Robust convex optimization. *Mathematics of Operations Research, 23*(4), 769–805. https://doi.org/10.1287/moor.23.4.769