




International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Joint Optimization of Capital Intensity and Workforce Flexibility in Distribution Centers Under Demand Uncertainty

Rohith Madhi Reddy¹, Nithin kumar Kanduri², Firoz Basha Syed³

¹ORCID ID - <https://orcid.org/0009-0004-5305-1335> Email: maviro1909@gmail.com

²ORCID ID - <https://orcid.org/0009-0008-4498-3156> Email: nkumarkanduri.gcc@gmail.com

³ORCID ID - <https://orcid.org/0009-0009-7166-6807> Email: Firozsyed974@gmail.com

Abstract

Distribution centers increasingly compete on the combined strength of two operating levers: the capital intensity of the facility, expressed through conveyor, sortation, goods-to-person, automated storage and retrieval, and robotic systems, and the flexibility of the workforce that runs it, expressed through the mix of permanent cross-trained associates and flexible or contingent labor. These two levers are typically planned by separate functions on separate timelines, even though both exist to absorb the same underlying source of risk: uncertain order volume. This paper develops a two-stage stochastic optimization model that jointly determines the capital intensity of a distribution center and the size of its core workforce before demand is realized, with flexible labor deployed as a recourse action once demand becomes known. A closed-form, newsvendor-type critical fractile characterizes the optimal core workforce for any given capital intensity, and the resulting expected cost function is minimized over capital intensity using a nested search procedure. Closed-form expressions for the expected recourse cost are obtained under a lognormal demand assumption, allowing exact evaluation without simulation. A synthetic numerical study, calibrated to representative ranges for large scale, multi-channel distribution operations, shows that jointly optimizing capital intensity and workforce flexibility reduces total expected annual cost by approximately eight percent relative to a capital investment decision made without access to workforce flexibility, and by more than fifty percent relative to a workforce only strategy that forgoes automation investment. Optimal capital intensity is shown to increase with demand uncertainty and with the wage premium commanded by flexible labor, and to decrease with the cost of capital. The results offer a quantitative basis for coordinating distribution center engineering and workforce planning decisions and identify the conditions under which each lever substitutes for, rather than complements, the other.

Keywords: distribution center design; capital intensity; workforce flexibility; demand uncertainty; two-stage stochastic programming; warehouse automation; contingent labor; capacity planning.

1. Introduction

Distribution centers occupy a position in modern supply chains that is both operationally central and financially demanding. A single large facility can process tens of millions of units per year, employ hundreds of associates, and require tens of millions of dollars in facility, material handling, and automation investment before it ever receives its first order. Two decisions determine how such a facility absorbs the volatility inherent in retail, e-commerce, and business to business demand. The first is the capital intensity of the operation: the extent to which conveyor, sortation, goods-to-person, automated storage and retrieval, autonomous mobile robots, and related technologies are used in place of manual material handling. The second is the flexibility of the workforce: the balance between a core group of permanent, cross-trained associates and a pool of flexible or contingent labor that can be added or removed as volume changes.

In most organizations these two decisions are made by different functions on different timelines. Facility and automation investment is typically evaluated by engineering or capital planning groups against a multi year horizon and a fixed hurdle rate, while workforce strategy is typically set by operations and human resources against a weekly or seasonal planning horizon. Each function optimizes its own lever while treating the other as a fixed input. This separation is understandable given the different skill sets and approval processes involved, but it is difficult to justify on economic grounds, because capital intensity and workforce flexibility are substitutes for the same underlying purpose: building operating capacity that can absorb uncertain demand

at the lowest expected cost. A facility with limited automation must lean more heavily on flexible labor to cover peak volume, while a highly automated facility reduces the labor hours required per unit processed and therefore changes the value of labor flexibility itself. Treating the two decisions independently risks a facility that is either over automated relative to the workforce strategy that supports it, or under automated relative to the labor market constraints it depends on.

The operations management literature offers well developed tools for each decision in isolation. The capacity and resource flexibility literature, beginning with the work of Fine and Freund (1990) and extended by Jordan and Graves (1995) and Van Mieghem (1998), establishes how investment in flexible manufacturing capacity should respond to demand correlation and uncertainty. A parallel literature on workforce and labor flexibility, including Sethi and Sethi (1990), Koste and Malhotra (1999), and the contingent labor contracting model of Milner and Pinker (2001), characterizes how a firm should size a core workforce and supplement it with contingent labor under uncertain requirements. A separate literature on warehouse and distribution center design, most comprehensively reviewed by de Koster, Le Duc, and Roodbergen (2007), addresses the engineering and operational tradeoffs involved in automation and material handling system selection. What is comparatively underdeveloped is a model that places the capital intensity decision and the workforce flexibility decision inside a single optimization problem, so that the substitution between them can be characterized formally rather than assumed.

This paper develops such a model. We formulate the joint decision as a two-stage stochastic program. In the first stage, before demand for a coming operating period is known, the distribution center operator chooses a capital intensity level, which determines the productivity of a labor hour, and a size for the core workforce, which is committed labor that cannot be adjusted without cost once demand is realized. In the second stage, demand is realized and the operator may deploy flexible or contingent labor, at a wage premium relative to core labor, up to a limit that reflects the practical capacity of temporary staffing arrangements. Any remaining shortfall between required and available labor capacity is met through overtime, expediting, or service failure, represented in the model as a penalty cost per labor hour of unmet capacity. This structure mirrors the operating reality described in supply chain engineering practice, where a distribution center's design basis specifies a core staffing plan alongside contingency provisions for peak and surge conditions.

The model yields several results with direct managerial relevance. First, for a fixed capital intensity, the optimal core workforce follows a closed-form critical fractile rule that generalizes the classical newsvendor solution to a setting with two sourcing tiers, core and flexible labor, and an explicit penalty for capacity that neither tier can cover. Second, the optimal capital intensity can be characterized through a nested search procedure that combines this closed-form inner solution with an outer search over the capital intensity variable, avoiding the need for large scale simulation or general purpose nonlinear solvers. Third, under a lognormal demand assumption, the expected cost of the recourse stage has a closed-form expression structurally analogous to a European option price, which makes the entire optimization problem exactly evaluable rather than approximated by Monte Carlo sampling. Fourth, a numerical study calibrated to representative parameter ranges for large scale, multi-channel distribution operations shows that the joint optimum reduces total expected annual cost by a material margin relative to policies that optimize capital intensity or workforce flexibility in isolation, and that the optimal capital intensity is sensitive to demand uncertainty, the wage premium associated with flexible labor, and the cost of capital in ways that are directionally intuitive but have not, to our knowledge, been jointly quantified in the distribution center context.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on capacity and resource flexibility, workforce and labor flexibility, and distribution center design, and identifies the gap this paper addresses. Section 3 describes the problem setting and states the modeling assumptions. Section 4 presents the formal two-stage stochastic optimization model, including the notation, the cost structure, and three propositions that characterize the optimal solution. Section 5 describes the nested solution procedure used to compute the joint optimum. Section 6 presents a synthetic numerical case study, including baseline results, a comparison against capital only and workforce only benchmarks, and sensitivity analysis with respect to demand uncertainty, the flexible labor wage premium, and the cost of capital. Section 7 discusses managerial implications for distribution center engineering and workforce planning practice. Section 8 discusses limitations of the model and directions for future research, and Section 9 concludes.

2. Literature Review

This paper draws on three streams of research that have largely developed in parallel: the operations management literature on capacity investment and resource flexibility under uncertainty, the literature on

workforce flexibility and contingent labor, and the literature on distribution center and warehouse design. A fourth, smaller stream on the labor market effects of automation and robotics provides additional context for the productivity assumptions used in the model. This section reviews each stream in turn and then summarizes the gap the present paper addresses.

2.1 Capacity Investment and Resource Flexibility Under Uncertainty

The foundational result in this literature is due to Fine and Freund (1990), who studied the optimal investment in product flexible manufacturing capacity when a firm faces uncertain demand across multiple products and must decide, before demand is realized, how much dedicated capacity to build for each product versus how much flexible capacity to build that can serve any product. Jordan and Graves (1995) extended this line of work with the influential finding that a small amount of strategically configured flexibility, in which each plant is linked to only a few products in a connected chain, captures most of the benefit of full flexibility. This chaining principle has become a standard reference point for flexibility design in both manufacturing and, as discussed below, workforce contexts.

Van Mieghem (1998) generalized the investment problem to a broader class of flexible resource technologies and characterized the optimal capacity portfolio using a marginal analysis that balances the cost of capacity against the expected marginal value of that capacity across demand states, a framework closely related to the critical fractile logic used in this paper. Van Mieghem (2003) later surveyed the capacity management and investment literature, organizing it around the central tradeoff between capacity cost, which is paid before demand is known, and the mismatch cost of having too much or too little capacity once demand is realized. Chod and Rudi (2005) extended the resource flexibility framework to settings with responsive pricing, showing that the value of flexibility depends on whether price can also adjust to realized demand. Bish and Wang (2004) examined how correlated demand across products affects the optimal level of flexible capacity investment, a consideration that is implicitly present in the present paper through the aggregate demand uncertainty faced by a single distribution center serving many stock keeping units and channels.

The stochastic programming methodology used in this paper, in which first stage decisions are made before uncertainty is resolved and second stage recourse decisions respond to the realized outcome, follows the standard two-stage formulation described by Birge and Louveaux (2011). Where the present model departs from typical capacity expansion formulations is in its use of a closed-form solution for the recourse stage, made possible by the lognormal demand assumption, rather than the sample average approximation or scenario based methods that are more commonly required when the recourse problem does not admit a closed-form. Bertsimas and Sim (2004) developed an alternative robust optimization approach to capacity and resource allocation under uncertainty that does not require a fully specified probability distribution; we adopt the stochastic programming approach in the main model because a fitted demand distribution is generally available in distribution center capacity planning from historical order data, but we note the robust approach as a natural extension in Section 8. The broader economic logic of investing under uncertainty, in which the option to delay or adjust investment has value, follows the real options framework of Dixit and Pindyck (1994); this framework is reflected in the present model's second stage recourse decision, which functions as a limited form of adjustment flexibility.

2.2 Workforce and Labor Flexibility

A parallel literature addresses flexibility on the labor side of the operation. Slack (1987) and Sethi and Sethi (1990) established an early taxonomy of manufacturing flexibility that explicitly includes labor and workforce flexibility as one of several flexibility dimensions alongside machine, routing, and volume flexibility. Koste and Malhotra (1999) refined this taxonomy, arguing that flexibility dimensions should be evaluated along consistent sub dimensions of range, mobility, and uniformity, a framework that has influenced subsequent empirical work on workforce agility. Upton (1995) argued, based on field research in manufacturing plants, that the flexibility realized by a facility depends less on the technology installed than on the range of conditions its people and processes can handle without excessive cost or delay, a finding directly relevant to the present paper's treatment of workforce flexibility as a decision variable that is jointly optimized with, rather than simply a consequence of, capital intensity.

Within the labor economics oriented side of operations management, Milner and Pinker (2001) modeled a firm's decision to size a permanent workforce and supplement it with contingent labor under uncertainty in both demand and the future availability of contingent workers, deriving a critical fractile solution for the optimal permanent staffing level. The core workforce sizing result in Section 4 of the present paper follows the same logic, extended to include a capital intensity choice that changes the labor hours required per unit of demand. Hopp, Tekin, and Van Oyen (2004) studied cross-training in serial production lines and showed that

a chaining strategy for cross-training workers, analogous to the chaining result of Jordan and Graves (1995) for manufacturing capacity, captures most of the benefit of full cross-training with a fraction of the training investment. Buke, Araz, and Fowler (2016) extended the cross-training literature to settings with imperfect training, in which cross-trained workers are not perfectly substitutable across tasks, a practical consideration relevant to distribution center associates who may be cross-trained across picking, packing, and replenishment functions but retain differential productivity across tasks. Together, this literature establishes that workforce flexibility, like capital flexibility, has diminishing but persistent value as the range and depth of flexibility increases, and that the design of flexibility, not simply its presence, determines the benefit captured.

2.3 Distribution Center Design, Warehouse Automation, and the Economics of Automation

The engineering and operational literature on distribution center and warehouse design provides the operational grounding for the capital intensity variable used in this paper. De Koster, Le Duc, and Roodbergen (2007) provide a comprehensive review of order picking system design, covering manual, mechanized, and automated picking strategies and the layout and control policies associated with each. Bartholdi and Hackman's widely used reference text on warehouse and distribution science documents the engineering relationships between storage system design, material flow, and labor productivity that underlie the productivity function used in Section 4 of this paper, in which higher capital intensity increases the units processed per labor hour subject to diminishing returns.

A smaller literature examines the labor market implications of automation and robotics more broadly. Acemoglu and Restrepo (2020) used variation in robot adoption across United States local labor markets to estimate the effect of robot deployment on employment and wages, finding a negative effect on employment in commuting zones most exposed to robot adoption. This finding is consistent with the productivity function used in this paper, in which higher capital intensity reduces the labor hours required per unit of throughput, and it motivates the framing of capital intensity and core workforce size as substitutes within the cost minimization problem rather than as independent choices.

2.4 Synthesis and Research Gap

Table 1 summarizes the literature reviewed above according to its primary focus, methodological approach, and treatment of the capital intensity and workforce flexibility decisions. Three observations follow. First, the capacity and resource flexibility literature is highly developed analytically but is framed primarily around manufacturing capacity investment rather than the combined capital and labor capacity decision faced by a distribution center operator. Second, the workforce flexibility literature independently develops closed-form staffing rules of the same critical fractile family used in the capacity literature, but generally treats the underlying capital or technology environment, and therefore labor productivity, as fixed. Third, the distribution center design literature addresses capital intensity and workforce requirements as engineering outputs of a design process rather than as jointly optimized decision variables under demand uncertainty. No study in the reviewed literature integrates capital intensity and workforce flexibility into a single stochastic optimization problem specific to distribution center capacity planning. The present paper addresses this gap by formulating and solving such a joint model, characterizing its solution analytically where possible, and quantifying the value of joint optimization relative to the siloed decision making that the literature review suggests is common in practice.

Table 1. Summary of related literature and positioning of the present paper.

Study	Primary focus	Method	Treatment of capital and labor
Fine and Freund (1990)	Flexible manufacturing capacity investment	Stochastic optimization	Capital only
Jordan and Graves (1995)	Manufacturing process flexibility, chaining	Analytical and simulation	Capital only
Van Mieghem (1998)	Investment in flexible resources	Marginal / stochastic optimization	Capital only
Van Mieghem (2003)	Capacity management and investment, review	Literature survey	Capital only

Study	Primary focus	Method	Treatment of capital and labor
Milner and Pinker (2001)	Contingent labor contracting	Two stage stochastic optimization, critical fractile	Labor only
Hopp, Tekin, and Van Oyen (2004)	Cross training, skill chaining	Queueing and simulation	Labor only
Buke, Araz, and Fowler (2016)	Cross training with imperfect schemes	Simulation and analytical	Labor only
de Koster, Le Duc, and Roodbergen (2007)	Warehouse order picking system design	Literature review	Engineering design; capital and labor as outputs, not jointly optimized
Acemoglu and Restrepo (2020)	Robots and local labor market employment	Empirical, econometric	Capital effect on labor, macro level
Present paper	Joint capital intensity and workforce flexibility in distribution centers	Two stage stochastic optimization, closed-form	Capital and labor jointly optimized under demand uncertainty

3. Problem Description and Assumptions

We consider a single distribution center that processes a random annual volume of order units, denoted D , across an operating period such as a fiscal year. Before the operating period begins, the facility operator commits to two decisions. The first is the capital intensity of the facility, denoted k and normalized to the interval from zero to one, where zero represents a fully manual, conventional racking and pick-to-pack operation and one represents a fully automated configuration incorporating automated storage and retrieval, autonomous mobile robots, and automated sortation. The second is the size of the core workforce, denoted n_c , measured in full-time equivalent associates who are cross-trained, permanently employed, and available for the full operating period regardless of realized demand.

Figure 1 summarizes the conceptual structure of the problem. Capital intensity and workforce flexibility jointly determine the distribution center's operating capacity, which must absorb realized demand at minimum expected cost.

Figure 1. Conceptual framework linking capital intensity, workforce flexibility, demand uncertainty, and total expected cost




Figure 1. Conceptual framework linking capital intensity, workforce flexibility, demand uncertainty, and total expected cost.

Figure 2 depicts the timeline of decisions. At time zero, before demand is realized, the operator commits to capital intensity and core workforce size. Between time zero and time one, demand uncertainty resolves. At time one, once demand is known, the operator may deploy flexible or contingent labor as a recourse action, and any capacity shortfall that flexible labor cannot cover is met through a shortfall penalty representing overtime, expediting, or service failure costs.




Figure 2. Two-stage decision timeline.

The model rests on the following assumptions, each grounded in standard distribution center operating practice.

Assumption 1, productivity and capital intensity. The number of order units that one labor hour can process, denoted θ of k , is increasing and concave in capital intensity k . This reflects the well documented pattern in warehouse engineering in which early automation investments, such as pick-to-light or basic conveyance, yield large productivity gains, while later investments toward full automation yield smaller incremental gains because a residual set of tasks, such as exception handling and quality checks, remain labor intensive regardless of the automation level installed.

Assumption 2, capital cost. The annualized cost of achieving capital intensity k , denoted C of k , is increasing and convex in k . Convexity reflects the increasing engineering complexity, systems integration risk, and construction cost associated with higher levels of automation, consistent with the capital expenditure patterns reported in industry sources on warehouse automation investment.

Assumption 3, demand distribution. Annual order volume D follows a lognormal distribution with mean μ_D and coefficient of variation CV , where the coefficient of variation captures the combined effect of seasonality, promotional activity, and channel mix uncertainty on order volume. The lognormal distribution is a standard choice in inventory and capacity planning applications because it restricts demand to positive values, exhibits right skew consistent with occasional high volume periods, and, as shown in Section 5, yields closed-form expressions for the expected cost of capacity shortfall.

Assumption 4, labor tiers and costs. Core labor is available at a fully loaded annual cost per full-time equivalent, denoted w_c , and provides H productive hours per year. Flexible or contingent labor is available only after demand is realized, at a cost per hour, denoted w_f , that exceeds the effective hourly cost of core labor, reflecting the wage premium, agency fees, and reduced productivity typically associated with temporary staffing in distribution operations. Flexible labor availability is capped at $F_{\max}$ of k hours, reflecting the practical limits of local contingent labor markets and the onboarding and training capacity of the facility. Demand for labor hours that neither core nor flexible labor can cover is met at a penalty cost per hour, denoted p , with p greater than w_f , representing overtime premiums, expedited outbound freight, or the cost of service failure.

Assumption 5, commitment timing. Capital intensity and core workforce size are first stage decisions, committed before demand is realized and not adjustable within the operating period without incurring costs outside the scope of this model, such as automation retrofit or workforce severance costs. Flexible labor deployment is a second stage recourse decision, made after demand is observed, consistent with the

operational reality that temporary staffing levels can be adjusted within an operating period while facility automation and permanent headcount cannot.

4. Model Formulation

Table 2 summarizes the notation used in the formal model.

Table 2. Notation.

Symbol	Description
k	Capital intensity of the distribution center, normalized to $[0,1]$
n_c	Core workforce, full-time equivalents (first stage decision)
$n_f(D)$	Flexible or contingent workforce deployed after demand D is realized (second stage decision)
D	Realized annual demand, in order units (random variable)
$\theta(k)$	Labor productivity at capital intensity k , in units per labor hour
$\theta_{\min}, \theta_{\max}$	Labor productivity at $k = 0$ and $k = 1$, respectively
λ	Shape parameter governing the productivity ramp with respect to k
$C(k)$	Annualized capital cost at capital intensity k
A	Annualized capital cost at full capital intensity ($k = 1$)
γ	Convexity parameter of the capital cost function
$L = D / \theta(k)$	Labor hours required to process realized demand D at capital intensity k
H	Productive hours provided by one core full-time equivalent per year
w_c	Fully loaded annual cost per core full-time equivalent
w_f	Cost per hour of flexible or contingent labor
$F_{\max}(k)$	Maximum available flexible labor hours, given capital intensity k
p	Penalty cost per hour of capacity shortfall beyond core and flexible labor
$S = L - H n_c$	Shortfall between required and core labor hour capacity
$Q(D, k, n_c)$	Second stage recourse cost, given realized demand D
$TC(k, n_c)$	Total expected annual cost (objective function)
μ_D, CV	Mean and coefficient of variation of annual demand D
μ_L, σ	Location and scale parameters of the lognormal distribution of L
F_L	Cumulative distribution function of the labor hour requirement L
$c(x; \mu, \sigma)$	Lognormal partial expectation, $E[\max(L - x, 0)]$, for $L \sim \text{lognormal}(\mu, \sigma)$

4.1 Productivity, Capital Cost, and Labor Requirement

Consistent with Assumption 1, we specify the productivity function as

$$\theta(k) = \theta_{\min} + (\theta_{\max} - \theta_{\min}) * (1 - \exp(-\lambda * k))$$

where $\theta_{\min}$ and $\theta_{\max}$ are the labor productivity, in units per labor hour, at zero and full capital intensity respectively, and λ governs how quickly productivity gains are realized as capital intensity increases. This functional form satisfies θ of zero equals $\theta_{\min}$, θ approaching $\theta_{\max}$ as k approaches one, and θ prime of k positive and decreasing, consistent with Assumption 1.

Consistent with Assumption 2, annualized capital cost is specified as

$$C(k) = A * k^{\gamma}, \gamma > 1$$

where A is the annualized cost of full capital intensity and γ controls the convexity of the cost function. Given realized annual demand D and capital intensity k , the labor hours required to process that demand are L equals D divided by θ of k . Because D is lognormal with parameters implied by μ_D and CV , and θ of k is deterministic given k , L is also lognormal, with the same coefficient of variation as D and a location parameter shifted by the log of θ of k .

4.2 Two-Stage Stochastic Program

The operator chooses k and n_c to minimize the sum of first stage costs and the expected second stage recourse cost:

$$\min_{k \in [0,1], n_c \geq 0} TC(k, n_c) = C(k) + w_c * n_c + E[Q(D, k, n_c)]$$

where the recourse cost, given realized labor hour requirement L equals D over θ of k , and letting S equal L minus H times n_c denote the shortfall between required and core labor hours, is

$$Q(D, k, n_c) = w_f * \min(\max(S, 0), F_{\max}(k)) + p * \max(S - F_{\max}(k), 0)$$

This recourse cost structure reflects a three tier labor sourcing hierarchy: core labor hours are used first, at cost w_c already sunk in the first stage; flexible labor hours are used next, up to the availability limit $F_{\max}$ of k , at cost w_f per hour; and any hours beyond the combined core and flexible capacity are met at the penalty cost p per hour. Because p exceeds w_f , which in turn exceeds the effective hourly cost of core labor, the cost structure creates an incentive to size core capacity and capital intensity so that the penalty tier is reached only in unlikely, high demand states.

4.3 Optimal Core Workforce for a Given Capital Intensity

We first characterize the optimal core workforce n_c for a fixed capital intensity k , ignoring for exposition the flexible labor capacity limit $F_{\max}$, which is reinstated in the numerical treatment in Section 5. Let F_L denote the cumulative distribution function of the labor hour requirement L equals D over θ of k .

Proposition 1. For fixed capital intensity k and unlimited flexible labor availability, the total expected cost TC of k comma n_c is convex in n_c , and the optimal core workforce n_c^* satisfies the critical fractile condition $F_L(H * n_c^*) = 1 - w_c / (w_f * H)$

Proof sketch. The expected recourse cost, E of $\max(S, 0)$, is a convex, decreasing function of the core labor hour capacity H times n_c , because it is the expectation of a convex function of a linear function of n_c . The first stage core labor cost w_c times n_c is linear in n_c . The sum of a linear and a convex function is convex, establishing convexity of TC in n_c . Differentiating the expected recourse cost with respect to n_c gives negative H times the survival function of L evaluated at H times n_c , because increasing core capacity by one hour reduces expected uncovered demand by exactly the probability that demand exceeds current core capacity. Setting the derivative of TC with respect to n_c to zero and solving for the survival function yields the stated condition.

This result is the distribution center capacity analogue of the classical newsvendor solution and of the contingent labor staffing rule derived by Milner and Pinker (2001): the core workforce should be sized so that the probability of a demand driven shortfall equals the ratio of the underage cost, here the core wage rate, to the cost of covering a marginal hour of shortfall with flexible labor. Two comparative statics follow directly from Proposition 1 and are used to interpret the sensitivity analysis in Section 6.

Proposition 2. The optimal core workforce n_c^* is nondecreasing in mean demand μ_D , and the optimal critical fractile $1 - w_c$ over w_f times H is nondecreasing in the flexible labor wage premium w_f over w_c . Consequently, holding capital intensity fixed, a higher flexible labor wage premium induces a larger optimal core workforce, and a lower flexible labor wage premium induces a smaller optimal core workforce, as the operator substitutes cheaper flexible labor for committed core capacity.

This proposition formalizes the intuitive substitution logic that motivates the joint model: because capital intensity k enters the labor hour requirement L through θ of k , any change in k that raises θ of k shifts the distribution of L to the left, which, by the critical fractile condition, generally reduces the number of core full-time equivalents required to reach a given service level, even though it need not reduce, and in the numerical results of Section 6 does not always reduce, the optimal capital intensity itself.

4.4 First-Order Condition for Capital Intensity

Let n_c^* of k denote the optimal core workforce characterized in Proposition 1 for a given k , and let TC^* of k equal TC of k comma n_c^* of k denote the resulting indirect expected cost function. By the envelope theorem, the derivative of TC^* with respect to k reflects only the direct effect of k on cost, holding n_c fixed at its optimal value, because the indirect effect operating through the optimal adjustment of n_c vanishes at the optimum.

Proposition 3. The optimal capital intensity k^* satisfies the first-order condition

$C'(k^*) = -d/dk \{ E[Q(D, k, n_c^*(k^*))] \}$ evaluated at $k = k^*$

which equates the marginal annualized capital cost of a small increase in capital intensity to the marginal reduction in expected recourse cost that the resulting productivity gain generates, holding the core workforce at its optimal level for that capital intensity. Because θ of k enters the labor hour requirement L multiplicatively through its reciprocal, and because C of k is convex while θ of k is concave, the right hand side of this condition is not generally available in closed-form once the flexible labor capacity limit $F_{\max}$ is active. Section 5 therefore solves for k star using a nested numerical procedure rather than a closed-form expression, while Proposition 3 provides the economic interpretation of the resulting solution: at the joint optimum, an additional unit of automation investment is justified only to the point where its annualized cost equals the expected labor cost savings it generates, inclusive of the reduced reliance on flexible and penalty rate labor that higher productivity makes possible.

5. Solution Methodology

Because the recourse cost function involves an expectation over a continuous demand distribution, the joint optimization problem in Section 4 could in principle be solved by sample average approximation, following the general stochastic programming methodology of Birge and Louveaux (2011), in which the expectation is replaced by an average over simulated demand draws and the resulting deterministic optimization problem is solved for increasing sample sizes until the solution stabilizes. Under the lognormal demand assumption maintained in this paper, however, the relevant expectations admit closed-form expressions, which allows the joint optimum to be computed exactly, without simulation error, and at negligible computational cost. This section describes the closed-form expressions and the nested search procedure used to compute the joint optimum.

5.1 Closed-Form Expected Shortfall

For a lognormal random variable L with location parameter μ_L and scale parameter σ , define the partial expectation function

$$c(x; \mu_L, \sigma) = E[\max(L - x, 0)] = E[L] * \Phi(d_1) - x * \Phi(d_2)$$

$$d_1 = (\mu_L + \sigma^2 - \ln(x)) / \sigma, \quad d_2 = (\mu_L - \ln(x)) / \sigma$$

where Φ denotes the standard normal cumulative distribution function and E of L equals $\exp$ of μ_L plus σ squared over two. This expression, derived in Appendix B, is structurally identical to the Black and Scholes expected payoff of a European call option on a lognormally distributed underlying asset, a parallel that is well known in the inventory and revenue management literature and that we exploit here purely as a computational device. Using this function, the expected recourse cost in the model of Section 4, inclusive of the flexible labor capacity limit, is obtained exactly as

$$E[Q] = w_f * c(H * n_c; \mu_L, \sigma) + (p - w_f) * c(H * n_c + F_{\max}(k); \mu_L, \sigma)$$

with μ_L equal to the location parameter of D minus the natural log of θ of k , and σ equal to the scale parameter implied by the demand coefficient of variation, which is invariant to the division of D by the deterministic quantity θ of k .

5.2 Nested Search Procedure

Figure 8. Nested solution algorithm
(outer capital-intensity search; inner closed-form workforce optimization)




Figure 8. Nested solution algorithm: outer capital-intensity search combined with inner closed-form workforce optimization.

Figure 8 summarizes the solution procedure. For any fixed capital intensity k , the inner problem of finding the cost minimizing core workforce n_c is a one dimensional convex minimization, established by Proposition 1, and is solved using a bounded scalar minimization routine applied directly to the closed-form total cost expression. The outer problem of finding the cost minimizing capital intensity k is solved by evaluating the inner optimal cost TC^* of k over a fine grid on the unit interval and identifying the minimizing value, a procedure equivalent to a golden section search given the empirically confirmed unimodality of TC^* of k illustrated in Figure 3 of Section 6. Because both the inner and outer problems are solved using closed-form or near closed-form expressions rather than simulation, the joint optimum, the associated sensitivity analysis, and the benchmark comparisons reported in Section 6 are computed to a precision limited only by the grid resolution used, which was set fine enough that further refinement did not change any reported result at the precision shown.

6. Numerical Study

This section applies the model and solution method of Sections 4 and 5 to a synthetic distribution center case. The parameters used are illustrative and calibrated to representative ranges reported in academic and industry sources for large scale, multi-channel distribution operations; they are not drawn from any specific facility, company, or client engagement, and the results should be interpreted as demonstrating the qualitative and approximate quantitative behavior of the model rather than as a forecast for any particular operation.

6.1 Case Description and Parameters

The synthetic distribution center processes a mean annual order volume of twenty five million units, with a baseline demand coefficient of variation of zero point two five, reflecting the seasonality and promotional variability typical of a multi-channel retail or consumer goods operation of this scale. Core associates provide two thousand productive hours per year at a fully loaded cost of twenty two dollars and fifty cents per hour, consistent with wage and benefit levels reported for distribution center associates in the United States. Flexible or contingent labor is available at thirty one dollars per hour, a premium of approximately thirty eight percent over the core wage rate, consistent with typical staffing agency markups, and is capped at forty percent of core labor hour capacity, reflecting practical limits on the availability and onboarding capacity of contingent labor in a given local labor market. Unmet capacity beyond the combined core and flexible tiers is penalized at fifty eight dollars per hour, reflecting a blend of overtime premiums and expedited freight or service failure costs. The productivity function ranges from twenty five units per labor hour at zero capital intensity to one hundred and ten units per labor hour at full capital intensity, with a shape parameter that concentrates most of the productivity gain in the lower to middle range of capital intensity, consistent with the pattern in which initial automation investments, such as conveyance and sortation, generate large productivity gains while later investments toward full automation generate smaller incremental gains. Annualized capital cost ranges from near zero at minimal capital intensity to six million dollars per year at full capital intensity, with a convexity parameter reflecting the increasing engineering and integration cost of higher automation levels. Table 3 summarizes these parameters.

Table 3. Case study parameters (illustrative and synthetic).

Parameter	Symbol	Baseline value
Mean annual demand	μ_D	25,000,000 units per year
Demand coefficient of variation	CV	0.25
Productive hours per core FTE	H	2,000 hours per year
Fully loaded core wage	w_c / H	\$22.50 per hour (\$45,000 per year)
Flexible labor wage	w_f	\$31.00 per hour
Flexible labor wage premium	$w_f / (w_c/H)$	1.38x core wage
Shortfall penalty rate	p	\$58.00 per hour
Flexible labor capacity limit	ϕ (fraction of core capacity)	0.40
Minimum productivity ($k = 0$)	$\theta_{\min}$	25 units per labor hour
Maximum productivity ($k = 1$)	$\theta_{\max}$	110 units per labor hour
Productivity ramp shape parameter	λ	3.0
Annualized capital cost at $k = 1$	A	\$6,000,000 per year
Capital cost convexity	γ	2.2

6.2 Baseline Results

Solving the joint optimization problem described in Sections 4 and 5 under the baseline parameters yields an optimal capital intensity of k^* equal to zero point four six, corresponding to a labor productivity of eighty

eight point six units per labor hour, an optimal core workforce of one hundred thirty point four full-time equivalents, and an annualized capital cost of approximately one point one million dollars. Total expected annual cost at the joint optimum is approximately eight point three one million dollars. This capital intensity level corresponds to a semi automated configuration incorporating conveyance, sortation, and goods-to-person picking technology rather than a fully manual or fully automated design, and the optimal core workforce is consistent with the staffing scale reported for distribution operations of comparable throughput in industry practice.

Figure 3 plots total expected annual cost as a function of capital intensity, with the core workforce re-optimized at each level of capital intensity according to Proposition 1. The cost function is convex in capital intensity across the range examined, confirming that the outer search procedure described in Section 5 locates a well defined interior minimum rather than a boundary solution. At low capital intensity, cost is dominated by the large core and flexible workforce required to process demand at low labor productivity. At high capital intensity, cost is dominated by the convex capital cost of achieving near full automation, which outweighs the labor savings that additional automation provides beyond the joint optimum.




Figure 3. Total expected annual cost as a function of capital intensity, with the core workforce re-optimized at each level of capital intensity.

Table 4 compares the joint optimum against three discrete automation tiers commonly used as reference points in distribution center design: a conventional, largely manual configuration; a semi automated configuration incorporating conveyor, sortation, and goods-to-person technology; and a highly automated configuration incorporating automated storage and retrieval, autonomous mobile robots, and robotic piece handling. The conventional configuration requires more than three hundred core full-time equivalents and results in a total expected annual cost more than double that of the joint optimum, driven primarily by core workforce cost. The highly automated configuration reduces workforce requirements substantially but incurs a capital cost that is not justified by the resulting labor savings, resulting in a total expected cost approximately thirty one percent higher than the joint optimum. The semi automated configuration, at a capital intensity close to but not identical to the joint optimum, performs nearly as well as the joint optimum itself, which is consistent with the relatively flat region of the cost curve around its minimum visible in Figure 3.

Table 4. Comparison of the joint optimum against discrete automation tiers.

Configuration	k	Productivity (units/labor-hr)	Annualized capex (USD)	Core FTE	Core labor cost (USD)	Total expected annual cost (USD)
Conventional (manual)	0.05	36.8	\$8,239	313.6	\$14,111,184	\$17,389,439
Semi-automated (conveyor/sortation/GTP)	0.45	88	\$1,035,664	131.3	\$5,909,805	\$8,314,961
Highly automated (AS/RS + AMR + robotics)	0.9	104.3	\$4,758,661	110.8	\$4,984,810	\$10,898,612
Joint-optimal k*	0.46	88.6	\$1,086,972	130.4	\$5,866,374	\$8,312,773

Table 4b decomposes the recourse-stage cost for each configuration into its expected flexible-labor and shortfall-penalty components, both of which shrink sharply once capital intensity moves beyond the conventional, manual configuration.

Table 4b. Expected flexible-labor and shortfall-penalty cost by configuration.

Configuration	E[flex labor hours/yr]	E[flex labor cost] (USD)	E[penalty hours/yr]	E[penalty cost] (USD)
Conventional (manual)	78,513	\$2,433,892	14,416	\$836,124
Semi-automated (conveyor/sortation/GTP)	32,881	\$1,019,321	6,037	\$350,171
Highly automated (AS/RS + AMR + robotics)	27,735	\$859,778	5,092	\$295,363
Joint-optimal k*	32,640	\$1,011,830	5,993	\$347,598




Figure 5. Expected annual cost decomposition by configuration, showing annualized capital cost, core workforce cost, expected flexible labor cost, and expected shortfall penalty cost.

6.3 Value of Joint Optimization

To quantify the benefit of jointly optimizing capital intensity and workforce flexibility, Table 5 compares the joint optimum against two benchmark policies that represent siloed decision making of the kind described in Section 1. The capital intensity only benchmark optimizes capital intensity and core workforce without access to flexible labor, representing an engineering led capacity planning process that does not incorporate contingent staffing into the design basis. The workforce flexibility only benchmark holds capital intensity near its minimum, representing an operations led staffing process that does not consider automation investment, while optimizing the core and flexible workforce mix. The capital intensity only benchmark results in a total expected annual cost approximately seven point eight percent higher than the joint optimum, because the absence of flexible labor forces a larger and more costly core workforce to protect against demand variability that flexible labor could otherwise absorb more cheaply. The workforce flexibility only benchmark results in a total expected annual cost more than fifty two percent higher than the joint optimum, because the absence of automation investment leaves labor productivity at its lowest level, requiring a very large core workforce regardless of how well that workforce is supplemented with flexible labor.

Table 5. Value of joint optimization relative to single-lever benchmark policies.

Policy	k	Core FTE	Total expected annual cost (USD)	Cost premium vs. joint optimum (%)
Joint optimization (proposed)	0.46	130.4	\$8,312,773	0.00%
Capital intensity only (no workforce flexibility)	0.475	145.2	\$9,013,317	7.77%
Workforce flexibility only (minimal automation)	0.05	313.6	\$17,389,439	52.20%

These results indicate that the cost of siloed decision making is asymmetric: underinvesting in workforce flexibility while optimizing capital intensity is materially less costly than underinvesting in capital intensity while optimizing workforce flexibility, at least under the baseline parameters examined here. This asymmetry follows from the convexity of the capital cost function relative to the near linear cost of core labor at low capital intensity, and it suggests that, of the two decisions, capital intensity is the more consequential lever when only one can be optimized well, though the joint optimum, which coordinates both, outperforms either single lever benchmark.

6.4 Sensitivity to Demand Uncertainty

Table 6 and Figure 4 report the joint optimum across a range of demand coefficients of variation from zero point one five to zero point four zero, spanning demand volatility levels from a relatively stable, mature retail channel to a highly volatile, promotion driven or new channel operation. As demand uncertainty increases, the optimal capital intensity increases modestly, from zero point four five zero at the lowest volatility examined to zero point four seven five at the highest, while total expected annual cost increases from approximately seven point eight seven million dollars to approximately nine point zero nine million dollars. The direction of this result, that higher demand uncertainty favors somewhat higher capital intensity, follows from the structure of the model: because higher capital intensity reduces the level of labor hours required for a given demand realization, it also reduces the absolute magnitude of the labor hour shortfall that a given percentage increase in demand volatility generates, partially offsetting the effect of increased uncertainty on expected recourse cost. The optimal core workforce is comparatively stable across the demand uncertainty range examined, indicating that in this model most of the adjustment to increased demand uncertainty occurs through the capital intensity decision and the resulting change in expected flexible labor and penalty costs, rather than through a change in committed core headcount.

Table 6. Sensitivity of the joint optimum to the demand coefficient of variation.

Demand CV	k*	theta*	Core FTE (n_c*)	Annualized capex (USD)	Total expected annual cost (USD)
0.15	0.45	88	131	\$1,035,664	\$7,874,203

Demand CV	k^*	θ^*	Core FTE (n_c^*)	Annualized capex (USD)	Total expected annual cost (USD)
0.2	0.455	88.3	130.4	\$1,061,149	\$8,077,619
0.25	0.46	88.6	130.4	\$1,086,972	\$8,312,773
0.3	0.465	88.9	130.4	\$1,113,135	\$8,565,297
0.4	0.475	89.6	130.1	\$1,166,480	\$9,088,952




Figure 4. Sensitivity of the joint optimum, capital intensity and core workforce, to the demand coefficient of variation.

6.5 Sensitivity to the Flexible Labor Wage Premium

Table 7 and Figure 6 report the joint optimum across a range of flexible labor wage premiums from fifteen percent to seventy percent above the core hourly wage, reflecting variation in local contingent labor market conditions. As the wage premium increases, both the optimal capital intensity and the optimal core workforce increase, from a capital intensity of zero point four five five and a core workforce of approximately one hundred twenty seven full-time equivalents at the lowest premium examined, to a capital intensity of zero point four six five and a core workforce of approximately one hundred thirty five full-time equivalents at the highest premium examined. This pattern is consistent with Proposition 2: as flexible labor becomes relatively more expensive, the model substitutes toward committed capacity, in the form of both a larger core workforce and, to a lesser extent, higher capital intensity, which reduces the volume of labor hours that must be sourced from either core or flexible tiers for a given demand realization.

Table 7. Sensitivity of the joint optimum to the flexible-labor wage premium.

Flexible-labor wage premium	w_f (USD/hr)	k^*	Core FTE (n_c^*)	Total expected annual cost (USD)
1.15	\$26	0.455	127.4	\$8,138,321
1.3	\$29	0.46	129.2	\$8,254,821
1.38	\$31	0.46	130.4	\$8,314,405
1.5	\$34	0.465	131.7	\$8,400,472

Flexible-labor wage premium	w _f (USD/hr)	k*	Core FTE (n _c *)	Total expected annual cost (USD)
1.7	\$38	0.465	134.7	\$8,534,894




Figure 6. Sensitivity of the joint optimum, capital intensity and core workforce, to the flexible labor wage premium.

6.6 Sensitivity to the Cost of Capital

Table 8 reports the joint optimum across a range of capital cost scale factors from seventy percent to one hundred thirty percent of the baseline annualized capital cost function, representing variation in equipment pricing, financing cost, or the discount rate used to annualize capital expenditure. As the cost of capital increases, the optimal capital intensity decreases, from zero point five one five at the lowest capital cost scale examined to zero point four two five at the highest, while the optimal core workforce increases correspondingly, from approximately one hundred twenty six to approximately one hundred thirty four full-time equivalents, as the model substitutes labor capacity for capital capacity when capital becomes relatively more expensive. This result is the expected direction implied by the first-order condition in Proposition 3 and confirms that the joint optimization framework correctly captures the substitution between capital and labor capacity as relative prices change.

Table 8. Sensitivity of the joint optimum to the cost of capital.

Capital cost scale	k*	Core FTE (n _c *)	Annualized capex (USD)	Total expected annual cost (USD)
0.7	0.515	125.7	\$975,495	\$7,945,462
0.85	0.485	128.1	\$1,038,011	\$8,139,979
1.00	0.46	130.4	\$1,086,972	\$8,312,773
1.15	0.44	132.3	\$1,133,561	\$8,468,819
1.3	0.425	133.9	\$1,187,272	\$8,611,416

Figure 7 summarizes the relative influence of the four parameters examined in Sections 6.4 through 6.6, together with mean annual demand, on total expected annual cost, using a tornado diagram constructed from

a plus or minus fifteen percent perturbation of mean demand and the parameter ranges used in the corresponding sensitivity tables. Mean annual demand is, unsurprisingly, the dominant driver of total expected cost, followed by demand uncertainty, the cost of capital, and the flexible labor wage premium, in that order. This ordering suggests that, in distribution center capacity planning practice, forecast accuracy for expected volume delivers more value than fine tuning the joint optimization model's secondary parameters, though the sensitivity of the joint optimum itself, as opposed to total cost, to demand uncertainty and the wage premium, documented in Sections 6.4 and 6.5, remains material for the capital intensity and workforce sizing decisions that are the focus of this paper.




Figure 7. Tornado diagram summarizing the sensitivity of total expected annual cost to mean demand, demand uncertainty, the cost of capital, and the flexible labor wage premium.

7. Managerial Implications

The results in Section 6 carry several implications for distribution center engineering and workforce planning practice. First, the magnitude of the cost premium associated with siloed decision making, approximately eight percent when capital intensity is optimized without access to workforce flexibility and more than fifty percent when workforce flexibility is optimized without capital investment, suggests that the organizational separation between facility engineering and workforce planning functions described in Section 1 carries a quantifiable economic cost. Coordinating these decisions, for example by including a representative labor cost and flexibility model in the automation investment business case, rather than treating labor cost as a fixed input to that case, has the potential to materially improve the total cost outcome of a distribution center design.

Second, the finding that optimal capital intensity increases with demand uncertainty runs somewhat counter to an intuition, sometimes expressed informally in capital planning discussions, that automation investment should be reserved for stable, predictable volume while uncertain or growing volume is better served by flexible labor because it avoids committing capital to a demand level that may not materialize. The model developed here suggests a more nuanced view: because higher capital intensity reduces the level of labor hours required per unit of demand, it also reduces the absolute scale of the labor capacity buffer, whether provided by flexible labor or by overtime and penalty rate labor, that a given degree of demand uncertainty requires. Automation and workforce flexibility are therefore better understood as complementary tools for managing demand uncertainty rather than as alternatives to be chosen between based on the perceived stability of the demand forecast.

Third, the sensitivity of the joint optimum to the flexible labor wage premium indicates that distribution center capital planning should incorporate current and projected local labor market conditions, not only current wage levels but the availability and cost premium of contingent labor, as an explicit input to the automation investment decision, rather than treating that decision as driven solely by facility throughput and service requirements. In labor markets where contingent staffing is scarce or expensive, the model indicates that a higher capital intensity investment is justified, all else equal, than in labor markets where flexible labor is abundant and inexpensive.

Fourth, the relatively flat region of the total cost curve around the joint optimum, visible in Figure 3 and reflected in the similar performance of the semi automated tier and the joint optimum in Table 4, suggests that precise identification of the exact optimal capital intensity is less important in practice than avoiding the two tails of the distribution, substantial underinvestment in automation or substantial overinvestment relative to the workforce strategy that supports it. This finding may be reassuring to practitioners, given that the parameters required to compute an exact joint optimum, including the shape of the demand distribution and the cost of capital, are themselves estimated with uncertainty in practice.

8. Discussion, Limitations, and Future Research

The model developed in this paper makes several simplifying assumptions that bound the scope of its conclusions and suggest directions for future research. First, the model considers a single, aggregate demand stream and a single labor pool, whereas real distribution centers typically process multiple product categories, order types, and fulfillment channels, each with distinct productivity, seasonality, and labor requirements. Extending the model to a multi product or multi-channel setting, in the spirit of the resource flexibility literature reviewed in Section 2.1, would allow the analysis to address the additional question of how capital intensity and workforce flexibility should be allocated across channels with correlated but not identical demand, rather than treated as a single aggregate decision.

Second, the model treats capital intensity as a continuous decision variable, which is a reasonable approximation for the purpose of characterizing the underlying tradeoff but abstracts from the discrete, lumpy nature of actual automation investment, in which a facility typically chooses among a small number of discrete technology packages rather than a continuous automation level. The discrete tier comparison in Table 4 partially addresses this limitation, but a fully discrete formulation, potentially solved as a mixed integer program, would better reflect the constrained choice set that a distribution center operator actually faces and could incorporate fixed costs associated with each discrete technology package that are not well approximated by the smooth, convex capital cost function used here.

Third, the model assumes that the demand distribution, including its coefficient of variation, is known and stationary over the operating period, whereas in practice demand forecasts are themselves uncertain and may be revised as new information arrives within the operating period. The robust optimization approach of Bertsimas and Sim (2004), noted in Section 2.1, offers an alternative to the stochastic programming formulation used here that does not require a fully specified demand distribution and could be applied to the same capital intensity and workforce flexibility decision to assess the sensitivity of the results to distributional misspecification, a natural direction for future research.

Fourth, the model treats flexible labor availability as an exogenous capacity limit that does not depend on the wage offered, whereas in practice the effective availability of contingent labor is likely to respond to the wage premium a facility is willing to pay, particularly in tight local labor markets. Incorporating an upward sloping flexible labor supply curve, rather than a fixed capacity limit, would allow the model to address the additional question of how the optimal wage premium itself should be set, alongside the capital intensity and core workforce decisions already addressed here.

Fifth, the numerical study in Section 6 uses a single synthetic case calibrated to representative parameter ranges for illustration. While the sensitivity analysis in Sections 6.4 through 6.6 demonstrates that the qualitative direction of the results is robust to substantial variation in key parameters, validating the model against operating data from an actual distribution center, where available, would strengthen confidence in the quantitative magnitude of the joint optimization benefit reported in Section 6.3, and represents a natural next step in a stream of applied research building on the framework developed here.

9. Conclusion

This paper has developed a two-stage stochastic optimization model that jointly determines the capital intensity and core workforce size of a distribution center under demand uncertainty, with flexible labor deployed as a second stage recourse action. The model yields a closed-form, critical fractile characterization of the optimal core workforce for a given capital intensity, a first-order condition characterizing the optimal capital intensity, and, under a lognormal demand assumption, exact closed-form expressions for the expected cost of the recourse stage that avoid the need for simulation. A synthetic numerical study calibrated to representative distribution center parameters shows that jointly optimizing capital intensity and workforce flexibility reduces total expected annual cost materially relative to policies that optimize either lever alone, and

that the optimal capital intensity responds in economically intuitive ways to demand uncertainty, the flexible labor wage premium, and the cost of capital. These findings support the view, motivated by supply chain engineering practice and developed formally in this paper, that capital intensity and workforce flexibility are complementary levers for managing demand uncertainty in distribution center operations and should be planned jointly rather than in the organizationally separated processes common in current practice. The model and solution method developed here provide a tractable starting point for such joint planning, and the limitations discussed in Section 8 outline several directions along which the framework can be extended to capture additional features of real distribution center operating environments.

Appendix A. Proof of Proposition 1

Fix capital intensity k and let L equal D divided by θ of k denote the resulting labor hour requirement, with cumulative distribution function F_L . Ignoring the flexible labor capacity limit for this derivation, the expected recourse cost as a function of core labor hour capacity x equals H times n_c is g of x equals w_f times the expectation of the maximum of L minus x and zero. This function is convex in x because it is the expectation, over the distribution of L , of a convex function of x for each realization of L , and an expectation of convex functions is convex. Differentiating under the integral sign, the derivative of g with respect to x equals negative w_f times the probability that L exceeds x , which is negative w_f times one minus F_L of x . Because total expected cost is the sum of the linear function w_c times n_c and the convex function g of H times n_c , total expected cost is convex in n_c . Setting the derivative of total expected cost with respect to n_c equal to zero gives w_c minus w_f times H times one minus F_L of H times n_c equals zero, which rearranges to the critical fractile condition stated in Proposition 1. Because total expected cost is convex in n_c , this first-order condition is both necessary and sufficient for a global minimum over n_c for the fixed capital intensity k .

Appendix B. Closed-Form Expected Shortfall for Lognormal Demand

Let L be lognormally distributed with the natural log of L normally distributed with mean μ_L and standard deviation σ . The partial expectation c of x equals the expectation of the maximum of L minus x and zero, for x greater than zero, is derived as follows. Write L equals the exponential of μ_L plus σ times Z , where Z is a standard normal random variable. Then

$$c(x) = \text{Integral over } z \text{ such that } \exp(\mu_L + \sigma z) > x \text{ of } [\exp(\mu_L + \sigma z) - x] * \phi(z) dz$$

where ϕ denotes the standard normal density. The region of integration is equivalent to z greater than d_2 , where d_2 equals the natural log of x minus μ_L , divided by σ . Splitting the integral into two terms, the first term, the integral of $\exp$ of μ_L plus σ times z times ϕ of z over the same region, evaluates, after completing the square in the exponent, to $\exp$ of μ_L plus σ squared over two times Φ of d_1 , where d_1 equals d_2 plus σ , and Φ denotes the standard normal cumulative distribution function. The second term, x times the integral of ϕ of z over the same region, evaluates to x times Φ of negative d_2 , which equals x times one minus Φ of d_2 . Combining terms and noting that $\exp$ of μ_L plus σ squared over two equals the expectation of L gives the closed-form expression used in Section 5, namely c of x equals the expectation of L times Φ of d_1 minus x times Φ of d_2 , with d_1 equal to μ_L plus σ squared minus the natural log of x , all divided by σ , and d_2 equal to d_1 minus σ . This expression is continuous, decreasing, and convex in x on the positive real line, properties used in the proof of Proposition 1 in Appendix A and confirmed numerically in the construction of Figure 3.

References

1. Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy* 128 (6), 2188 to 2244.
2. Arrow, K.J., Chenery, H.B., Minhas, B.S., Solow, R.M., 1961. Capital labor substitution and economic efficiency. *Review of Economics and Statistics* 43 (3), 225 to 250.
3. Bartholdi, J.J., Hackman, S.T., 2019. Warehouse and Distribution Science, Release 0.98. Georgia Institute of Technology, Atlanta, Georgia.
4. Bertsimas, D., Sim, M., 2004. The price of robustness. *Operations Research* 52 (1), 35 to 53.
5. Birge, J.R., Louveaux, F., 2011. Introduction to Stochastic Programming, second ed. Springer Series in Operations Research and Financial Engineering. Springer, New York.
6. Bish, E.K., Wang, Q., 2004. Optimal investment strategies for flexible resources, considering pricing and correlated demands. *Operations Research* 52 (6), 954 to 964.
7. Buke, B., Araz, O.M., Fowler, J.W., 2016. Cross training with imperfect training schemes. *Production and Operations Management* 25 (9), 1611 to 1629.

8. Chod, J., Rudi, N., 2005. Resource flexibility with responsive pricing. *Operations Research* 53 (3), 532 to 548.
9. Chopra, S., Meindl, P., 2019. *Supply Chain Management: Strategy, Planning, and Operation*, seventh ed. Pearson, New York.
10. de Koster, R., Le Duc, T., Roodbergen, K.J., 2007. Design and control of warehouse order picking: A literature review. *European Journal of Operational Research* 182 (2), 481 to 501.
11. Dixit, A.K., Pindyck, R.S., 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, New Jersey.
12. Fine, C.H., Freund, R.M., 1990. Optimal investment in product flexible manufacturing capacity. *Management Science* 36 (4), 449 to 466.
13. Hopp, W.J., Tekin, E., Van Oyen, M.P., 2004. Benefits of skill chaining in serial production lines with cross trained workers. *Management Science* 50 (1), 83 to 98.
14. Jordan, W.C., Graves, S.C., 1995. Principles on the benefits of manufacturing process flexibility. *Management Science* 41 (4), 577 to 594.
15. Koste, L.L., Malhotra, M.K., 1999. A theoretical framework for analyzing the dimensions of manufacturing flexibility. *Journal of Operations Management* 18 (1), 75 to 93.
16. Milner, J.M., Pinker, E.J., 2001. Contingent labor contracting under demand and supply uncertainty. *Management Science* 47 (8), 1046 to 1062.
17. Sethi, A.K., Sethi, S.P., 1990. Flexibility in manufacturing: A survey. *International Journal of Flexible Manufacturing Systems* 2 (4), 289 to 328.
18. Silver, E.A., Pyke, D.F., Thomas, D.J., 2017. *Inventory and Production Management in Supply Chains*, fourth ed. CRC Press, Boca Raton, Florida.
19. Skinner, W., 1974. The focused factory. *Harvard Business Review* 52 (3), 113 to 121.
20. Slack, N., 1987. The flexibility of manufacturing systems. *International Journal of Operations and Production Management* 7 (4), 35 to 45.
21. Upton, D.M., 1995. What really makes factories flexible. *Harvard Business Review* 73 (4), 74 to 84.
22. Van Mieghem, J.A., 1998. Investment strategies for flexible resources. *Management Science* 44 (8), 1071 to 1078.
23. Van Mieghem, J.A., 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing and Service Operations Management* 5 (4), 269 to 302.
24. Zipkin, P.H., 2000. *Foundations of Inventory Management*. Irwin McGraw Hill, Boston.