Business Analytics for Dual-Channel SaaS Pricing under Demand, Supply, and Cybersecurity Risk

Xinze Xu[1], Haining Wang[1, *], Jianfeng Wu[1]

[1] Business School, Beijing International Studies University, Beijing, China 100024

* Email: wanghaining2023@163.com (Corresponding Author)

Received 10 December 2025 · Revised 24 January 2026 · Accepted 10 March 2026 · Published 30 June 2026

Abstract

Software-as-a-service providers increasingly combine direct subscriptions with sales through resellers, yet their pricing decisions must also account for uncertain demand, limited cloud capacity, and cybersecurity losses. This study develops a prescriptive business analytics model for a provider and a coordinated reseller that sell a common subscription service through two channels. The research uses a transparent four-state illustrative scenario set rather than proprietary observations. Price-dependent, cross-channel demand, a capacity-related service-credit obligation, and state-dependent cyber incident probabilities are integrated in a single expected-profit objective. Security expenditure reduces incident probabilities exponentially and incurs a convex cost. Conditions under which the optimization problem is concave are derived, and the resulting constrained problem is solved with sequential quadratic programming. For the specified illustrative parameters, the best coordinated policy sets direct and reseller prices at USD 113.17 and USD 113.48 per subscription period, reserves 241.21 capacity units, and chooses a dimensionless security effort of 0.789. Expected profit is USD 17,624.47, compared with USD 17,306.71 when security effort is fixed at zero. Changing the joint likelihood of high demand and disrupted supply while holding marginal probabilities constant alters the recommended capacity and the magnitude of expected service credits. These results show how correlated operating and cyber risks can be translated into pricing and investment decisions, but they are scenario-dependent and do not constitute empirical validation. The model provides a reproducible baseline for testing contract, forecasting, and resilience decisions with organizational data.

Keywords: subscription services; channel coordination; cloud capacity; service credits; security investment; scenario analysis; prescriptive optimization

Business Analytics for Dual-Channel SaaS Pricing under Demand, Supply, and Cybersecurity Risk

1. Introduction

Software-as-a-service (SaaS) businesses must coordinate commercial promises with the technical systems that make those promises credible. A subscription price is not only a revenue number: it determines the demand that must be supported, the resources reserved in advance, the customer relationships exposed to interruptions, and the potential scale of a security incident. Unlike a single transaction in which fulfillment ends at purchase, a subscription creates a continuing expectation of availability, assistance, data protection, and renewal. The economics of subscription and usage-based service therefore require the pricing function to work with operations and security teams rather than operating independently. The relevant unit of analysis is the combination of a commercial offer and the capabilities needed to deliver it (Katzan, 2009; Li & Kumar, 2022; Ojala, 2016).

Many providers reach customers through more than one route. A direct channel may sell through the provider's website or enterprise sales team, while an indirect channel may use a reseller, marketplace partner, or systems integrator. The two channels can differ in support cost, customer acquisition effort, contract administration, and the value of additional services. They also compete for some of the same buyers. A change to one price can attract new customers, shift customers from the other channel, or change the work required of a partner. The resulting decision is related to the broader literature on dual-channel distribution, but the capacity and service obligations of cloud software make a simple transfer of physical-goods assumptions incomplete (Chiang et al., 2003; Dan et al., 2012; Zhang & Wang, 2018).

The business problem becomes more difficult when three uncertainties interact. First, demand depends on price, buyer heterogeneity, seasonality, and the effectiveness of channel partners. Second, cloud resources available to the provider can be reduced by infrastructure failure, provider-side throttling, or failures in a critical dependency. Third, an incident can produce remediation expenses, contractual liabilities, and customer-related losses. These uncertainties need not occur independently. An organization under heavy demand may face an infrastructure shortfall while an external disruption also raises the probability of a cyber incident. Decision makers who evaluate each exposure using separate spreadsheets can miss the cost of their joint occurrence (Bryce, 2019; Feng et al., 2022; Liu & Wei, 2022).

The operational question addressed here is: how should a coordinated dual-channel SaaS business select prices, capacity, and security effort when demand, supply availability, and incident exposure are uncertain and potentially dependent? The question is prescriptive. It asks what a decision maker should choose under a stated objective and assumptions, not whether a historical intervention caused a particular observed result. The analysis distinguishes the risk of failing to deliver a contracted service from lost sales: the base model assumes that the subscription fee is collected and that an explicit service credit is paid for shortfall. This convention makes the accounting transparent and permits the use of a concave formulation, although it does not represent every cloud contract.

The literature provides useful building blocks but often treats them in different settings. SaaS studies examine tariff design, licensing, operations, and adoption; distribution studies analyze channel conflict and coordination; and security research investigates incident probabilities, incentives, and breach costs (Hausken, 2006; Wu et al., 2024; Xin, 2020; Zhang et al., 2020). Our contribution is not the discovery of each individual mechanism. It is their integration into one bounded, reproducible decision model with shared capacity, two linked prices, a conditional incident probability, and an auditable service-credit mechanism. A finite scenario representation makes dependencies visible and enables the same framework to be populated with forecasts or expert assessments when suitable data become available.

The paper makes three contributions. First, it expresses cross-channel demand and infrastructure shortfall in the same decision model, explicitly separating economic channel costs from transfers between a provider and a reseller. Second, it links the conditional cyber probability and loss exposure to the demand–supply states, derives a sufficient condition for global concavity, and identifies the capacity and security investment conditions. Third, it applies the model to a fully specified synthetic example, compares risk and mitigation configurations, and tests changes in both security cost and the joint distribution of demand and disruption. The numerical values are demonstrations of model behavior, not measured market evidence.

The remainder of the article reviews the relevant literature, defines the decision setting, develops the optimization model, reports the numerical and sensitivity results, and discusses implementation and research limitations. Because the model is evaluated with hypothetical inputs, the findings should be interpreted as conditional decision insights. They are not forecasts of the prices, cyber probabilities, or financial performance of a real SaaS provider.

2. Literature Review

2.1 Subscription pricing and customer choice

A SaaS offer is typically delivered as a continuing service, although the charging mechanism can be flat-rate, usage-based, or a combination of an access fee and usage charges. The appropriate tariff depends on demand heterogeneity, usage patterns, customer value, and the costs of serving different users. The research on subscription-based information goods shows that pricing arrangements affect both revenue and customer selection, while software licensing research emphasizes that uncertainty about a product's value can influence a buyer's initial choice (Fishburn & Odlyzko, 1999; Xin, 2020; Yao et al., 2012). A fixed per-period subscription in our model is a deliberate simplification, not a claim that it is universally optimal.

Bundling and tier design can produce very different usage and profitability patterns even when customers face the same advertised price. A provider may also adjust its pricing approach as its cloud business evolves from a perpetual-license model toward recurring service delivery (Ojala, 2016; Wu et al., 2024). In the present setting, both channels offer the same underlying service and differ only in demand intercepts, direct economic costs, and prices. This restricted comparison isolates the consequences of interacting risks. Introducing quality differentiation, bundles, or customer-specific usage would be valuable, but doing so without corresponding data would obscure the interpretation of the numerical example.

Customer adoption depends on perceived value and the organizational consequences of moving applications into a cloud environment. Research on cloud adoption provides a basis for considering organizational fit and implementation risks alongside financial savings (Khajeh-Hosseini et al., 2012; Safari et al., 2015). These determinants can be represented in a richer demand intercept or in customer-segment models. Here, the channel intercepts are given parameters; they are not estimated from survey responses, web traffic, or actual subscription records. The policy recommendation should therefore be read as an optimization conditional on these intercepts and the assumed linear price responses.

2.2 Dual channels and coordination

Dual-channel research has demonstrated that a direct channel can change the incentives of an intermediary and the competitive position of the overall supply system (Chiang et al., 2003). Services delivered by a reseller may create value even when a customer ultimately purchases directly, and customer channel preferences may alter pricing and coordination outcomes (Dan et al., 2012; Xu et al., 2023). Other models incorporate partner sales effort, service value, environmental attributes, or social responsibility in the pricing problem (Li et al., 2020; Ranjan & Jha, 2019; Zhang & Wang, 2018). These studies illustrate the range of possible contracting environments rather than establishing one universal channel structure.

Our model considers a coordinated system in which the provider and reseller optimize total economic profit over the two customer-facing prices. Coordination may arise because the reseller acts as an agent, because both prices are agreed within a lawful commercial arrangement, or because the system-level solution is used as a planning benchmark. No particular contract form is assumed to implement it. A wholesale transfer or commission paid by the provider to the partner is a redistribution within the system and is excluded from total profit, whereas the partner's real selling and support expenditure remains in the reseller-channel unit cost. This distinction prevents commissions from being counted as a system cost merely because one party pays the other.

The coordinated solution is not a claim that an independently acting reseller will choose the same price. If a partner maximizes its own margin, a decentralized response and the provider's subsequent decision must be modeled separately; equilibrium, participation, and legal constraints may all matter. A coordinated benchmark remains useful because it identifies the surplus attainable under its assumptions and makes the potential cost of misaligned incentives measurable once an appropriate decentralized model is added. The present paper does not estimate that coordination gap or prescribe a minimum resale price.

2.3 Supply and service continuity

Supply risk in cloud services differs from a shortage of physical inventory, but both affect the feasibility of promised service. Research on service agreements and resource provisioning emphasizes that availability, workload, and allocated resources must be considered when determining cloud commitments (Ghandour et al., 2024; Wang et al., 2017; Yuan et al., 2023). Pricing and provisioning also evolve as providers learn about workloads and service costs (Chen et al., 2023; Gera & Xia, 2011). A provider may contract for a fixed reservation, supplement it with elastic resources, or accept exposure to service credits. The relevant economic comparison is between the cost of additional reliable capacity and the expected cost of a shortfall.

Operational resilience research similarly distinguishes a disturbance from its consequences. The same outage can have different financial effects depending on preparation, restoration capabilities, and the demand present at the time (Ferdous et al., 2025; Scala & Lindsay, 2021; Singh & Singh, 2019). Analytics capability can help organizations identify and respond to supply disruptions, but access to data does not itself guarantee that operational or commercial decisions are coordinated (Liu & Wei, 2022). Our model treats availability as a state-specific fraction of reserved capacity. It does not model recovery time, cloud-region topology, backup facilities, or the engineering process that generates an outage.

Service-level agreements can translate an operational shortfall into a financial claim. Cloud research discusses the relationship between security, performance, compensation, and contractual obligations (Chen et al., 2022; Luna et al., 2017). We use an illustrative credit per unit of contracted demand not supported by available capacity. This aggregate representation is suitable when contract terms are converted into an expected seat-equivalent credit, but it is not a substitute for reviewing actual service-level agreements. In particular, the credit should be redefined if payment is lost rather than refunded, if liability is capped, or if the service agreement measures downtime rather than unsupported subscribers.

2.4 Cybersecurity as an economic decision

Security is not simply an external cost to be added after the price has been chosen. The size of a customer base and the contractual obligations associated with it can affect expected incident losses, while security investment can change the probability of a damaging event. Economic models of security investment show that the shape of the breach-probability function and the decision maker's preferences can materially alter the investment recommendation (Gao et al., 2015; Hausken, 2006; Huang et al., 2008). We use an exponential probability response and expected profit for tractability; other probability functions and risk preferences are possible and may lead to different results.

Cloud-specific research examines how user interdependence, service-provider competition, customization, and security exposure interact (Arce, 2020; Feng et al., 2022; Zhang et al., 2020). Security governance and the management of software security requirements provide additional operational context for interpreting an abstract effort variable (Bryce, 2019; Ramachandran, 2016). The model does not claim that all cyber events originate from the same source or that the same control reduces every type of incident equally. Its variable z represents an aggregate program whose cost and effect would need to be calibrated for a defined incident class and time horizon.

Privacy and breach consequences can extend beyond the immediate cost of restoring a system. Research considers privacy risks in cloud computing, the response of financial markets to breach announcements, and variation in breach risk across firms (Goel & Shawky, 2009; Svantesson & Clarke, 2010; Wang & Ngai, 2022). These findings motivate including fixed and subscriber-related components in our loss function, but they do not supply the numerical parameters used here. A real assessment should avoid double counting: if lost revenue is already embedded in a per-subscriber loss estimate, it must not be included again as a separate incident-related sales reduction. The base model also does not equate a breach with a supply outage; the two can be correlated through the scenario probabilities without one being assumed to cause the other.

2.5 Analytics, identification, and decision quality

Forecasting research has developed statistical and machine-learning approaches to predict demand and adjust forecasts for business decisions (Nanty et al., 2025; Seyedan & Mafakheri, 2020; Wibowo, 2024). Hybrid forecasting methods can incorporate price information and temporal structure, but better predictive accuracy alone does not establish the causal effect of changing a price (Cantalupo & Silverio, 2026; Ruan et al., 2022). Historical prices are often set in response to anticipated demand, channel promotions, and capacity limits. A regression that ignores such endogeneity may suggest an elasticity that produces a poor intervention, even if it tracks past observations reasonably well (Bates & Kim, 2024; Miller & Alberini, 2016).

Causal inference tools can help identify the effect of a pricing intervention when the design and assumptions support such interpretation (Guelman & Guillén, 2014; Rose & Rizopoulos, 2019; Zhao & Liu, 2023). For a dual-channel business, identification must include cross-channel substitution and any changes in the composition of customers. An observed increase in direct sales following a price reduction can reflect new demand, migration away from the reseller, or both. Network models of product substitution provide a related perspective on the interdependence of prices and demands (Chen & Lim, 2023). Our synthetic linear slopes are assumed rather than estimated, and no causal effect is inferred from the numerical results.

Scenario-based optimization is one way to convert predictions into a decision, but its recommendation depends on the specified probability distribution. Distributionally robust methods consider a set of plausible distributions instead of treating one empirical distribution as certain (Liu et al., 2023; Luo & Mehrotra, 2020; Rahimian & Mehrotra, 2022). Related methods use alternative risk measures or moments to represent adverse outcomes (Liu et al., 2017; Long & Qi, 2014). We discuss these extensions but do not claim to solve a distributionally robust or risk-averse model. The base case maximizes expected profit under a single, explicitly listed finite distribution, which keeps the accounting and sensitivity calculations reproducible.

3. Methodology

3.1 Research design and system boundary

This is a methodological study with a synthetic numerical experiment. There are no observations from a particular business, no respondents, and no claim of an empirical field test. One provider supplies a subscription service through a direct channel, denoted d, and a coordinated reseller channel, denoted r. Decisions are made at the start of a single subscription period. Both customer-facing prices, reserved capacity,

and security effort are chosen before the demand–availability state is revealed. All scenarios use the same time horizon; profit and all monetary inputs are expressed in USD per period.

Demand is interpreted as the number of contracted subscription-equivalent units in the period. For simplicity, one contracted unit requires one capacity-equivalent unit. All contracts pay their price regardless of a temporary shortfall, and the provider then pays the explicit service credit. The system bears a direct operating cost c_i for each contracted unit in channel i, a capacity reservation cost k for every reserved unit, and a security program cost. These costs cover distinct activities. The fee paid by the provider to the reseller is a transfer and is not deducted again from coordinated profit. This accounting represents total provider–reseller economic profit, not the separate profit of each organization.

The scenario is s = 1, …, S, with probability $\pi_s \geq 0$ and sum equal to one. In each state, the demand shock and the fraction of reserved capacity that is available are jointly specified. A cybersecurity event is characterized by its conditional probability given state s and security effort z. The event affects the financial loss term but does not automatically create an additional capacity shortage; any association between the event and infrastructure availability is represented by conditioning. Timing, boundaries, and the principal variables are summarized in Table 1.

Table 1. Principal Notation

Symbol	Definition
d, r; s	Direct and reseller channels; operating state
p_i, $D_{i,s}$	Subscription price and contracted demand in channel i and state s
a_i, b_i, g	Demand intercept, own-price coefficient, and cross-price coefficient
σ_i, ε_s	Channel shock scale and state-specific demand shock
π_s, u_s	State probability and available fraction of reserved capacity
K, H_s, k, η	Capacity reserved; shortfall; reservation cost; unit service credit
c_i	Economic operating cost per contracted unit in channel i
ρ_s, z, α	Baseline incident probability, security effort, and cost coefficient
F, ℓ	Fixed and per-contract loss conditional on a cyber event
B, S(p)	Demand coefficient matrix and probability-weighted unmitigated loss

Note. Monetary terms are in USD per subscription period; capacity and demand are measured in equivalent units.

3.2 Demand and cross-channel substitution

Let $p = (p_d, p_r)$ be the per-period customer-facing price vector. Demand in state s follows the linear system in Equation 1, in which a_i is the channel intercept, b_i is the own-price coefficient, g is a symmetric cross-price coefficient, σ_i scales the common demand shock, and ε_s is its state-specific realization. The common shock keeps the example concise; a richer application can add channel-specific shocks and correlated residuals.

$$D_{d,s} = a_d - b_d p_d + g p_r + \sigma_d \varepsilon_s$$
$$D_{r,s} = a_r - b_r p_r + g p_d + \sigma_r \varepsilon_s \quad (1)$$

The own-price coefficients are positive and g is nonnegative. A positive g means that raising one channel's price increases the other channel's demand, holding the other variables fixed. It does not identify whether the additional demand is entirely customer migration or an increase in market participation. B is the symmetric demand matrix with diagonal entries b_d and b_r and off-diagonal entries −g. We require $b_d b_r > g^2$, which makes B positive definite and ensures diminishing revenue in the two price directions. The chosen bounds also keep

both state-specific demands nonnegative; otherwise truncation or an alternative demand model would be required.

Prices are bounded by a pre-specified commercial range, $30 \leq p_i \leq 140$ in the numerical example. Capacity K is bounded between zero and 600 units, and security effort z between zero and two. These limits make the optimization domain explicit and represent the range over which the assumed relationships are considered credible. They are not legal or market-wide limits. The unconstrained linear equations should not be extrapolated indefinitely, particularly where predicted demand would become negative or price sensitivity changes materially.

3.3 Capacity shortfall and service credits

A state-specific availability factor u_s lies between zero and one. The usable capacity in state s is u_sK; the contracted volume is $D_{d,s} + D_{r,s}$. Equation 2 defines the seat-equivalent shortfall H_s as the positive part of their difference. A coefficient η converts each shortfall unit into a service credit. For example, if 10 units of demand are unsupported and the credit is USD 30 per unit, the modeled shortfall cost is USD 300. The model treats this payment as a financial credit rather than a reduction in the number of contracts sold.

$$H_s = max\{ D_{d,s} + D_{r,s} - u_s K, 0 \} \quad (2)$$

This definition is deliberately aggregate. It does not decide which subscribers or channels receive the available capacity, and it assumes a common credit rate. If channel-specific service guarantees exist, H_s could be replaced by separate channel shortfalls with an allocation rule. If users consume different resource quantities, demand should first be converted into a common capacity unit using a documented workload model. The per-unit credit must also be consistent with the period and with the contractual cap on refunds. Under the baseline parameters the credit does not exceed the minimum permitted subscription price; in every reported optimum the unit credit is below the corresponding channel prices.

The scenario model represents the combined effect of high or low demand and normal or impaired supply. Holding demand and disruption marginals fixed while changing their joint probabilities isolates the effect of dependence. When high demand and low availability coincide more frequently, shortfalls may become more expensive even if the separate demand forecast and outage rate do not change. A manager looking only at the marginal averages would not see this shift in tail exposure. The numerical experiment reports the joint probabilities, not merely the separate averages.

3.4 Security exposure and investment

For a given state s, let ρ_s be the probability of a specified loss-producing cyber event when no additional security effort is undertaken. The security policy reduces that conditional probability according to Equation 3. The scalar z is dimensionless and is not a dollar amount; it is a normalized effort level. Its monetary cost is $\alpha z^2/2$, where α is a positive cost coefficient. The exponential function is an explicit illustrative assumption with diminishing absolute returns and positive residual risk for finite z.

$$Pr(cyber\ event \mid s, z) = \rho_s\, exp(-z) \quad (3)$$

If an incident occurs, the system incurs a fixed loss F plus ℓ times the total contracted volume in that state. F may represent incident response and restoration; ℓ may represent contractual or customer-related loss per contracted unit. The components must be defined so that the same loss is not charged twice. We assume that the parameters F and ℓ do not themselves depend on z; a security program that changes severity rather than probability would require a different specification. The conditional probability ρ_s can vary with supply availability, but the model does not infer causation from that association.

The expected cyber loss is $\exp(-z)$ times the probability-weighted sum of $\rho_s[F + \ell(D_{d,s} + D_{r,s})]$. If high cyber exposure tends to coincide with high demand, expected per-subscriber loss depends on the joint distribution rather than the product of unconditional averages. In an application, the incident category, observation

window, exposure base, and treatment of near misses should be fixed before estimating ρ_s. Because severe incidents are rare, estimates may be uncertain; sensitivity ranges and external evidence may be more informative than a single historical average.

3.5 Integrated objective and optimality

The coordinated objective is the expected subscription margin less expected service credits, expected cyber loss, the capacity reservation charge, and the security program cost. Equation 4 presents the objective. The optimization selects p_d, p_r, K, and z within their stated bounds; all scenarios share these decisions. The objective does not include a separate risk-aversion penalty, so a higher variance of profit has no independent cost unless it changes expected credits or incident losses. Introducing conditional value-at-risk would be a distinct model rather than an alternative interpretation of this objective.

$$max\ \Pi\ =\ \Sigma_s\ \pi_s\ [\ (p_d\ -\ c_d)D_{d,s}\ +\ (p_r\ -\ c_r)D_{r,s}\ -\ \eta H_s$$
$$-\ \rho_s\ exp(-z)\{F\ +\ \ell(D_{d,s}\ +\ D_{r,s})\}\ ]\ -\ kK\ -\ \alpha z^2/2 \quad (4)$$

The credit term creates a piecewise-linear kink whenever usable capacity equals demand in a state. A numerically convenient equivalent uses nonnegative variables t_s with $t_s \geq D_{d,s} + D_{r,s} - u_s K$. Because η is positive, the optimum sets t_s equal to H_s for each scenario with positive probability. This epigraph representation has linear feasibility constraints and avoids differentiating the positive-part function directly. It also makes it straightforward to inspect the slack and verify that credits are being computed from the same demand and availability assumptions as the revenue component.

The basic model admits a useful global concavity condition. Write $S(p) = \Sigma_s\ \pi_s\rho_s[F + \ell(D_{d,s} + D_{r,s})]$ and $v = \nabla_p S$. With symmetric B positive definite, nonnegative demand, $z \geq 0$, $S(p) \geq 0$, and $\alpha > v^T(2B)^{-1}v$, the objective is concave jointly in p, K, and z. The quadratic subscription-margin term has Hessian $-2B$; the negative shortfall penalty is concave because the positive-part function is convex; and the security term satisfies the Schur-complement condition under the stated bound. The feasible set is convex. Thus, a feasible solution satisfying the first-order conditions is globally optimal for this formulation. The sufficient condition is conservative and is not asserted to be necessary.

For the illustrative parameters, the eigenvalues of B are approximately 1.084 and 1.716. The security cross-effect bound $v^T(2B)^{-1}v$ is 1.651, compared with $\alpha = 500$ in the base case. These quantities support the stated concavity certificate over the bounded domain. The condition also holds for the two security-cost sensitivities reported later. No claim is made that the same certificate applies if demand is truncated, cyber losses are nonlinear in demand, a discrete contract choice is added, or the effort effect changes form.

When an interior security choice is optimal, the marginal cost of security equals the marginal expected loss avoided, as in Equation 5. At a boundary, the corresponding one-sided inequality replaces the equality. The security decision is therefore related to both the conditional incident probabilities and the price-dependent number of contracted subscriptions. The relationship is economic rather than technical: z has no universal physical meaning, and two organizations can require very different controls and budgets to achieve the same proportional reduction in incident probability.

$$\alpha z\ =\ exp(-z)\ \Sigma_s\ \pi_s\ \rho_s\ [F\ +\ \ell(D_{d,s}\ +\ D_{r,s})] \quad (5)$$

The analogous capacity condition depends on states in which demand exceeds available capacity. Away from a kink, the marginal expected reduction in credits from one additional reserved unit is $\eta\Sigma_s\ \pi_s u_s$ over shortage states. At a kink, Equation 6 gives the subgradient interval. Capacity is worthwhile up to the point at which its reservation cost is balanced by this expected credit reduction. A discrete scenario model can place the optimum exactly at a demand-to-availability threshold; therefore capacity can change sharply when the credit rate or joint probabilities cross a threshold, even if underlying parameters change smoothly.

$$\eta \Sigma_{H_s > 0} \pi_s u_s \leq k \leq \eta \Sigma_{H_s \geq 0} \pi_s u_s \quad (6)$$

3.6 Data design, inputs, and solution procedure

The experiment uses four illustrative demand–availability states. Their probabilities sum to one and are not derived from observed cloud service data. The low-demand and high-demand shocks are −1.1 and 0.9, respectively; because their marginal probabilities are 0.45 and 0.55, the expected shock is exactly zero. Normal availability is 1.0 and impaired availability is 0.7. Conditional incident probabilities before security effort are 0.08 in normal-supply states and 0.25 in impaired-supply states. These values are scenario inputs, not measured security incident rates. Table 2 presents the full distribution and the remaining parameters.

Table 2. Illustrative Scenario Distribution and Model Inputs

Component	Definition	Value
State 1	Normal supply; low demand: (π, ε, u, ρ)	(0.40, −1.1, 1.0, 0.08)
State 2	Normal supply; high demand: (π, ε, u, ρ)	(0.35, 0.9, 1.0, 0.08)
State 3	Impaired supply; low demand: (π, ε, u, ρ)	(0.05, −1.1, 0.7, 0.25)
State 4	Impaired supply; high demand: (π, ε, u, ρ)	(0.20, 0.9, 0.7, 0.25)
Demand	(a_d, a_r); (b_d, b_r); g	(250, 200); (1.5, 1.3); 0.3
Demand shocks	(σ_d, σ_r)	(25, 20)
Operating costs	(c_d, c_r); capacity cost k	USD (12, 18); USD 5
Service credit	η per unit shortfall	USD 30
Cyber loss	F; ℓ per contract	USD 5,000; USD 10
Security cost	α in $\alpha z^2/2$	USD 500
Decision bounds	(p_d, p_r); K; z	[30, 140]; [0, 600]; [0, 2]

Note. All values are hypothetical. The state tuple is probability, shock, availability, and baseline incident probability, respectively.

The four-state probabilities imply a 25% marginal probability of impaired supply and a 55% marginal probability of high demand. In the base case the probability of simultaneous high demand and impaired supply is 20%. The high-demand probability conditional on impaired supply is consequently 80%, compared with its 55% unconditional rate. The model also specifies a higher cyber probability during impaired supply. These conditional relationships create dependence without requiring the cyber event to be the cause of the outage or the high demand. A scenario is simply a joint description of the assumed operating environment.

We solve the epigraph form with Python 3, NumPy 2.3.5, and SciPy 1.17.0 using the sequential least-squares programming implementation in scipy.optimize.minimize. The solver uses an objective tolerance of 10^{-9}, at most 2,000 iterations, and three different starting points. A solution is retained only when all linear shortfall inequalities are satisfied within 10^{-5}. The concavity condition supplies the global interpretation for the stated formulation; multiple starts and numerical residuals are checks on implementation, not a substitute for that mathematical condition. The complete runnable specification and starting points are provided in Supplement S1. Values are rounded only for display.

The scenario comparisons use the same underlying commercial parameters. The deterministic benchmark sets the shock standardizers to zero, sets all availability factors to one, and removes cyber loss. The demand-only case restores the zero-mean demand shock, while the next case also restores impaired supply. Two full-risk cases differ in whether security effort is fixed at zero or optimized. The order of these comparisons is a reporting convention: allocating the total change to individual risks would generally depend on the order in

which risks are introduced, and the reported increments should not be interpreted as a unique causal decomposition.

3.7 Validation and sensitivity design

The objective and every parameter needed to recalculate it are specified in the paper. Validation consists of verifying that the probabilities sum to one, the demand shock has the claimed mean, all prices and capacity decisions satisfy their bounds, computed demands are nonnegative, the epigraph shortfalls equal the positive-part expression, and the profit decomposition sums to the objective. For the full-risk optimized case, the reported solution lies at a capacity kink, so the capacity subgradient rather than a single ordinary derivative is used to check the first-order condition.

Sensitivity analysis changes one group of assumptions at a time and reoptimizes all decisions. Security cost is varied from $\alpha = 250$ to 1,000. The service-credit rate is varied from $\eta = 20$ to 45. A dependence experiment changes the probability of the high-demand and impaired-supply state from 0.05 to 0.125 and 0.20 while keeping both marginals fixed. For each chosen joint probability j, the state probabilities become $(0.20 + j, 0.55 - j, 0.25 - j, j)$. This construction is feasible for $0 \le j \le 0.25$ and avoids confusing a change in dependence with a change in the overall outage or high-demand rate.

The analysis does not use sampling-based inference. There are no estimated coefficients, standard errors, confidence intervals, or p-values because the parameters are deliberately assigned and the four scenario probabilities are treated as known within each model run. Sensitivity ranges reflect alternative assumptions, not statistical uncertainty intervals. If the model is calibrated from empirical data, uncertainty in elasticities, incident frequencies, contract costs, and state probabilities should be propagated through the entire optimization rather than appended only to the final profit estimate.

4. Results

4.1 Coordinated prices and capacity across scenarios

Table 3 reports the optimized decisions and expected profit for the five configurations. The deterministic no-risk benchmark produces a direct price of USD 111.99, a reseller price of USD 112.31, capacity of 203.30 units, and expected profit of USD 18,814.12. With zero-mean demand variation but no supply or cyber risk, the optimal prices remain essentially the same, while capacity rises to 243.80 units to cover the high-demand state. Expected profit falls to USD 18,611.62 because the additional reservation incurs cost. This result follows from the specific discrete distribution and credit–capacity cost ratio; it does not imply that uncertainty always leaves prices unchanged.

Table 3. Optimized Decisions across Illustrative Risk Cases

Case	Direct price	Reseller price	Capacity	Effort z	Expected profit
No risk	111.99	112.31	203.30	0	18,814.12
Demand only	111.99	112.31	243.80	0	18,611.62
Demand and supply	112.89	113.21	241.82	0	18,174.57
All risks; z = 0	113.51	113.82	240.47	0	17,306.71
All risks; z optimized	113.17	113.48	241.21	0.789	17,624.47

Note. Prices and profit are USD per period; capacity is in equivalent units. Figures are synthetic and rounded.

Adding impaired supply changes the recommended prices to USD 112.89 and USD 113.21, and capacity becomes 241.82 units. Expected profit declines to USD 18,174.57. Although demand and supply uncertainty often motivate extra capacity, the comparison with the demand-only case shows a slightly lower optimum

here. The business also raises prices, which reduces contracted demand, and the optimal capacity remains at the normal-supply high-demand threshold. The result illustrates why capacity and price should be interpreted jointly rather than assuming a monotonic response for every parameter configuration.

With all three risks present and security effort fixed at zero, the solution yields expected profit of USD 17,306.71. Allowing security effort to be optimized increases expected profit to USD 17,624.47, an improvement of USD 317.75, or about 1.84%, relative to the unprotected case. The protected policy sets prices of USD 113.17 and USD 113.48, reserves 241.21 capacity units, and selects $z = 0.789$. Relative to the same risky system without additional protection, the optimal prices are slightly lower and contracted demand is slightly higher because the expected per-customer cyber exposure has fallen. This is a local comparison within the stated model, not a general rule for all security investments.

The incremental comparison from demand-only to demand-plus-supply reduces expected profit by USD 437.06. Adding unmitigated cyber exposure to the latter reduces the optimum by a further USD 867.85. Because the decisions are reoptimized after each change, these differences combine the direct financial effect of the added risk with the response of prices and capacity. Changing the sequence in which supply and cyber risks are introduced would change the allocation of their combined effect. None of the differences is an observed firm loss or an estimate of a causal intervention.

4.2 Profit decomposition and first-order checks

The protected configuration yields an expected subscription margin before capacity, credits, and cyber costs of USD 19,814.61. Expected service credits are USD 434.17; expected cyber-event loss is USD 394.39; reserved capacity costs USD 1,206.04; and the security program costs USD 155.54. These figures sum, subject to rounding, to USD 17,624.47. The mean modeled shortfall is 14.47 capacity-equivalent units per period. It is an expected shortfall in the finite scenario distribution, not a forecast measured in a production system.

Before the additional security program, the probability-weighted incident probability is 0.1225. At $z = 0.789$, the multiplier $\exp(-z)$ is approximately 0.454 and the weighted probability becomes 0.0557, a reduction of approximately 54.6% under the assumed exponential response. The expected incident loss does not fall by exactly the same proportion across the reoptimized cases because the prices and subscriber volumes also change. In the protected optimum, the marginal security cost αz and the marginal expected loss avoided $\exp(-z)S(p)$ both equal approximately USD 394.39 per unit of z, consistent with Equation 5.

In the base solution, the high-demand normal-supply state lies at the capacity boundary, while the high-demand impaired-supply state has a positive shortfall. With $\eta = 30$, the expected marginal credit reduction from the strictly short states is $30 \times 0.20 \times 0.70 =$ USD 4.20 per extra unit of reserved capacity. Including the boundary state raises the upper end of the subgradient interval to $30 \times (0.20 \times 0.70 + 0.35) =$ USD 14.70. The capacity cost of USD 5 falls within this interval. This verifies the capacity optimality condition without pretending that the kink has a unique derivative.

The epigraph variables match the nonnegative shortfall expression within 10^{-5}, and the numerical first-order residuals for the two prices and the security effort are below 0.001 in absolute terms at the unrounded solution. The reported values are rounded, so a reader using the rounded prices may observe slightly larger residuals. Concavity, bounded feasibility, and the subgradient check support a global optimum for this specific model and parameter set; they do not validate the realism of the synthetic parameters or the adequacy of the linear demand model.

4.3 Security cost and service-credit sensitivity

Table 4 reports selected sensitivities. When the security cost coefficient falls to $\alpha = 250$, optimal effort increases to 1.126 and expected profit rises to USD 17,734.43. When α increases to 1,000, effort falls to 0.517 and expected profit is USD 17,523.21. These comparisons hold the baseline cyber probabilities and

loss parameters fixed while changing only the cost of obtaining the proportional risk reduction. The monetary expenditure $\alpha z^2/2$ need not fall proportionally with α because the chosen level of effort also changes.

Table 4. Sensitivity of the Optimized Policy

Sensitivity	Direct price	Reseller price	Capacity	Effort z	Expected profit
Base case	113.17	113.48	241.21	0.789	17,624.47
Security cost $\alpha = 250$	113.09	113.41	241.38	1.126	17,734.43
Security cost $\alpha = 1{,}000$	113.26	113.57	241.02	0.517	17,523.21
Credit $\eta = 20$	112.87	113.18	241.87	0.789	17,769.39
Credit $\eta = 45$	113.34	113.66	344.04	0.789	17,541.83
Joint adverse state $j = 0.05$	112.50	112.81	242.69	0.778	17,961.58
Joint adverse state $j = 0.125$	112.84	113.15	241.95	0.783	17,792.76
Joint adverse state $j = 0.20$	113.17	113.48	241.21	0.789	17,624.47

Note. Only the indicated assumption changes; all decisions are reoptimized. j is the joint probability of high demand and impaired supply; marginals remain fixed.

The service-credit sensitivity highlights the discrete nature of the capacity decision. With $\eta = 20$, the optimal capacity is 241.87 units and expected profit is USD 17,769.39. With $\eta = 45$, capacity rises to 344.04 units and modeled shortfalls in all four states disappear, at a lower expected profit of USD 17,541.83. In the latter case, additional reservation cost is justified by the more expensive credits. Profit falls when the penalty becomes larger because the decision maker is facing less favorable contractual terms; the increased capacity is a response to those terms, not evidence that increasing capacity itself reduces profit in every setting.

The sharp change from approximately 241 to 344 units should not be mistaken for a precise real-world tipping point. It is produced by four discrete states and a linear unit credit. A continuous workload distribution, alternative capacity products, nonlinear penalties, or outages with partial duration could smooth or move the threshold. For managerial use, such sensitivity is a prompt to examine the contract and the cloud reservation menu, not a directive to purchase the exact modeled capacity. Reporting the state distribution is essential for understanding why the recommendation changes.

4.4 Dependence between demand and disruption

The dependence experiment holds the 55% probability of high demand and the 25% probability of impaired supply constant. At a joint probability of 0.05 for high demand and impairment, the optimized expected profit is USD 17,961.58 and expected service credits are USD 109.21. At an independent joint probability of 0.125, expected profit is USD 17,792.76 and credits are USD 272.19. At the base-case joint probability of 0.20, expected profit is USD 17,624.47 and credits are USD 434.17. Across these scenarios the independent case has the same marginal forecasts as the other two but a different alignment of adverse conditions.

Expected profit decreases by USD 337.11, about 1.88%, between the 0.05 and 0.20 joint-probability cases. The optimized direct price increases from USD 112.50 to USD 113.17, while capacity decreases from 242.69 to 241.21 units. In this numerical range the price response reduces demand and the binding threshold moves, even though expected credit exposure increases. The direction of capacity adjustment is conditional on the same scenario structure and cost ratios; it should not be generalized to every setting in which dependencies become stronger.

The dependency effect also reaches the cyber-loss component. Because incident probabilities are larger in impaired-supply states and the event loss includes a subscriber-related amount, changing the alignment of

high demand with impairment changes the probability-weighted exposure even with the same unconditional incident probability before security effort. The security choices differ slightly across the three cases. This mechanism differs from claiming that an outage causes an attack: the model represents statistical or scenario dependence, and identifying a causal mechanism would require different evidence and a different research design.

5. Discussion

5.1 What the model adds to SaaS pricing research

The central finding is that a coordinated price is partly an operations decision. In the model, a price changes both the revenue per subscription and the volume exposed to capacity constraints and potential cyber losses. This link complements research that connects SaaS tariff selection to usage and operations (Li & Kumar, 2022; Wu et al., 2024) and research that studies capacity and pricing as joint decisions (Chen et al., 2023; Gera & Xia, 2011). The formal result is narrower than a general prescription: it applies to fixed-period subscription fees, the specified linear demand system, a common capacity pool, and expected profit.

The dual-channel formulation adds a second commercial consequence. Raising the direct price reduces direct demand but also changes reseller demand through g; the same relationship works in the opposite direction. A channel-by-channel calculation that ignores substitution can therefore misstate the value of either a price change or capacity reservation. This does not mean the provider can necessarily set both prices in practice. Where a reseller has independent pricing authority, the coordinated optimum should be treated as a benchmark, and the actual decision requires the reseller's feasible response and the terms of the relationship (Chiang et al., 2003; Dan et al., 2012; Xu et al., 2023).

The numerical study also shows that the strongest statement is about the value of a modeling discipline rather than a universal price direction. Price and capacity can move in opposite directions as the assumed distribution changes. A larger expected shortfall can coexist with lower chosen capacity because the price response reduces demand. The direction of any one variable depends on service-credit rates, the distribution of states, capacity prices, and customer substitution. The model makes those assumptions explicit enough to be challenged and revised.

5.2 Translating the framework into a practical analytics pipeline

A deployment would begin by defining the decision horizon and the unit of service. A monthly subscription, an annual enterprise contract, and a burstable usage plan have different exposure periods and contractual commitments. The provider and partner should then agree which customer-facing prices are truly controllable, which costs are resource costs, and which payments are transfers. If a reseller's support team is part of the economic system, its labor and infrastructure costs belong in c_r; a commission paid to the same partner should not be counted again as an external cost. Consistent accounting is a prerequisite for an interpretable optimization.

The next step is assembling demand information. A practical dataset might join historical prices, contract dates, channel attribution, lost opportunities, promotions, product versions, geography, account size, and renewal status. Supply information could include reserved capacity, utilization, performance, dependency incidents, duration, and credits paid. Security information could include event definitions, exposure windows, incident severity, and independent control assessments. These examples specify what an organization could collect; no such data were used in this paper. Data should be linked at a level that supports the decision while limiting access to sensitive or identifiable information.

A third step is estimating demand responses with attention to selection. The firm should distinguish a price that was set because demand was already expected to rise from demand that would change if the firm actually altered the price. Where lawful and commercially acceptable, randomized or quasi-experimental price

changes can strengthen identification. Otherwise, documented confounders, time effects, credible instruments, or sensitivity analyses may be needed. Predictive cross-validation and causal validity address different questions: an accurate forecast based on historical pricing policies can still be unreliable when used to evaluate a new price policy (Guelman & Guillén, 2014; Miller & Alberini, 2016; Zhao & Liu, 2023).

A fourth step is building joint rather than separate risk scenarios. Demand residuals, availability events, and cyber incidents should be aligned on the same period and exposure definition. The organization should investigate reporting changes, censored incidents, missing outages, and rare-event uncertainty before using simple historical frequencies. When the data do not support stable joint estimates, expert elicitation or scenario ranges can be used, with the source and uncertainty documented. Such ranges can then feed stress tests or an explicitly formulated distributionally robust model (Liu et al., 2023; Luo & Mehrotra, 2020; Rahimian & Mehrotra, 2022). These procedures remain proposed applications, not results established by the present synthetic experiment.

5.3 Cybersecurity, service governance, and contractual design

The security effort variable should be translated into a defined set of controls before a budget recommendation is made. For example, a team could compare identity controls, vulnerability management, recovery capabilities, and monitoring programs for a specified incident class. Each program may change probability, severity, detection time, or recovery time differently. Combining those effects into a single exponential parameter is convenient for a first model but can conceal important differences between controls. The relevant comparison is an incremental, evidence-based change from an existing baseline, not an assumption that security starts at zero.

Cyber costs should be coordinated with legal, finance, and engineering estimates. A service credit, an incident response invoice, a regulatory obligation, and customer attrition may overlap or occur in different periods. In the base model, a capacity credit is paid when contracted demand exceeds available capacity, whereas the cyber term covers a separately defined incident loss. If an actual cyber incident causes downtime and produces the same contractual credit, the organization must model the joint cash flows explicitly rather than adding both costs uncritically. The same care is needed when a provider and a reseller share responsibility for incident response (Gao et al., 2015; Luna et al., 2017; Svantesson & Clarke, 2010).

Service-level terms can change the economically desirable capacity threshold. A contract with a cap, grace period, service measurement window, or tiered refund does not map automatically to a linear per-seat penalty. Similarly, common definitions of cloud uptime may refer to time availability, not the count of unsupported customers. The model's seat-equivalent shortfall can be calibrated only after the contract is translated into equivalent financial exposure. A careful implementation would simulate the actual agreement on historical or synthetic workload traces and then compare its credits with the simplified function (Chen et al., 2022; Ghandour et al., 2024).

5.4 Partner incentives and implementation limits

A coordinated system optimum does not ensure that each party benefits from adopting it. The provider may bear most security and cloud costs, while the reseller may bear acquisition and support costs or depend on a commission. Before the recommendation can be implemented, both parties would need a feasible agreement on responsibilities, data access, compensation, and risk sharing. A contract could be evaluated by adding the provider's and reseller's separate profit equations and checking that each prefers participation to its outside option. Those equations are not estimated or solved here.

The distinction between a system objective and a contract is especially important for data sharing. A reseller might possess valuable information about its customer pipeline but be reluctant to reveal it; the provider might possess sensitive incident or infrastructure data. The value of additional information depends on how it changes the decision, not simply on whether a forecasting metric improves. A governance agreement should

specify data access, permitted use, validation, retention, and audit rights. More accurate information can reduce uncertainty, but incentives and commercial constraints determine whether it is shared and acted upon.

Decision governance should also protect against the mechanical use of uncertain optimization results. A manager can require a range of feasible prices, an explanation of the forecast and its limitations, independent review of major changes, and a record of the inputs used. Alerts should identify when demand, outages, or incidents fall outside the scenario set. A recalculation schedule and an escalation route are more useful than treating one optimized number as permanently correct. Pilot deployments can compare realized margins and service outcomes with a baseline policy, while recognizing that the baseline must be specified before the evaluation.

5.5 Limitations and future research

The first limitation is the absence of empirical calibration. The 53 cited studies provide conceptual and methodological context, but none is a source for the invented numerical values in Tables 2–4. Therefore the analysis cannot establish that the proposed price, probability reduction, or profit improvement applies to a particular provider. An empirical extension should use observed contracts and relevant operational records, describe the sampling frame and cleaning steps, and test decisions out of sample. If confidentiality prevents public release of transaction-level data, synthetic examples and an auditable model can still be shared alongside an appropriate restricted-access procedure.

Second, the linear, symmetric cross-price structure simplifies behavior. In practice, a direct price may have a different effect on reseller demand than a reseller price has on direct demand, and the response may vary by segment or price range. Demand may also be censored when systems cannot serve all interested buyers. Allowing a nonsymmetric demand matrix requires re-examining the curvature of the revenue function and the sufficient conditions for concavity. If the commercial relationship involves tiers, bundles, switching costs, or churn, the objective and the definition of demand should be extended before making quantitative recommendations.

Third, the service-credit representation collects the full subscription fee despite the modeled capacity shortfall. That assumption may be inappropriate if outages cause cancellations, refunds above the modeled credit, reputational loss, lost sales, or a reduction in future renewals. A multi-period model could link current service quality and security incidents to later adoption and churn. Such an extension would also require consistent incident timing and a clear account of how customers update their expectations. Double counting between future churn and immediate cyber losses should be avoided.

Fourth, the scenario distribution is small and assumed known. Real cyber events are infrequent and their occurrence probabilities can be hard to estimate; dependencies may change as infrastructure, attackers, and demand patterns evolve. A mean-profit objective can also obscure unacceptable downside exposure. Future work can incorporate conditional value-at-risk, ambiguity sets, stress scenarios, risk constraints, and alternative cyber probability functions (Hausken, 2006; Huang et al., 2008; Liu et al., 2017; Liu et al., 2023; Long & Qi, 2014). The resulting model should be labeled as risk-averse or distributionally robust only when the corresponding measure or ambiguity set is actually specified and solved.

Fifth, the model does not solve a decentralized game, optimize a wholesale contract, or represent competing SaaS providers. A future extension could compare a coordinated benchmark with a provider–reseller game in which each party has a separate objective and the reseller chooses its own price. Such a model would require a transparent timeline, participation constraints, and evidence about the feasible contract space. It could also assess whether a proposed agreement shares the cost of security and service credits in a way that both parties accept. These are research questions rather than conclusions of the present analysis.

Finally, the security response is a reduced-form approximation. More effort does not always affect every incident class smoothly, and controls can interact or fail together. Incident severity may depend on the

number and type of affected users rather than the number of subscriptions alone. A richer operational model could distinguish prevention, detection, response, and recovery and include shared infrastructure dependencies. Its increased detail must be supported by credible data; a more elaborate equation is not automatically a more reliable decision tool.

6. Conclusion

This study develops a coordinated two-channel SaaS pricing model that links customer demand, capacity availability, service credits, and cybersecurity investment. For a clearly specified class of assumptions, the model is concave and admits interpretable conditions for capacity and security decisions. A four-state synthetic illustration shows how the alignment of high demand with impaired supply affects expected credits and the optimized commercial policy, and how security investment changes expected profit under a specified probability response.

The contribution is a reproducible analytical benchmark and a way to organize the information needed for a practical decision. The numerical values are not estimates for any actual company. Before operational or publication use, the assumptions, input calibration, contracts, author declarations, and interpretations require review by the responsible researchers and decision makers.

Acknowledgement

We acknowledge the researchers whose published work informed this study. Any remaining errors are our responsibility.

Author Contributions

Conceptualization, X.Z.X. and H.N.W.; Methodology, X.Z.X. and H.N.W.; Formal Analysis, X.Z.X. and H.N.W.; Investigation, X.Z.X. and J.F.W.; Data Curation, X.Z.X. and J.F.W.; Software, X.Z.X. and J.F.W.; Validation, X.Z.X. and H.N.W. and J.F.W.; Resources, J.F.W.; Visualization, X.Z.X. and J.F.W.; Writing – Original Draft, X.Z.X.; Writing – Review & Editing, X.Z.X. and H.N.W. and J.F.W.; Supervision, H.N.W.; Project Administration, X.Z.X. and H.N.W. We have read and approved the final manuscript and take responsibility for the work.

Funding

This research received no external funding.

Conflicts of Interest

We declare no conflicts of interest.

Ethical Approval and Consent

The analysis uses hypothetical inputs and does not involve human participants, personal data, or animals.

Data Availability Statement

All scenario probabilities, model parameters, objective functions, and solution settings used in the illustrative numerical experiment are reported in the manuscript. The accompanying Supplement S1 contains executable Python code for the reported scenarios. No proprietary or empirical dataset was used.

Declaration of Generative AI Use

We used OpenAI ChatGPT (GPT-6) to assist with literature discovery, drafting, analytical implementation, and document preparation. We take full responsibility for the accuracy of the references, calculations, and final manuscript.

References

Arce, D. G. (2020). Cybersecurity and platform competition in the cloud. Computers & Security, 93, 101774. https://doi.org/10.1016/j.cose.2020.101774

Bates, M., & Kim, S. (2024). Estimating the price elasticity of gasoline demand in correlated random coefficient models with endogeneity. Journal of Applied Econometrics, 39(4), 679–696. https://doi.org/10.1002/jae.3042

Bryce, C. N. (2019). Security governance as a service on the cloud. Journal of Cloud Computing, 8(1), 23. https://doi.org/10.1186/s13677-019-0148-5

Cantalupo, M., & Silverio, F. (2026). Price-dependent demand forecasting of consumer packaged goods: A hybrid Prophet-LSTM approach. Journal of Business Analytics, 1–26. https://doi.org/10.1080/2573234x.2026.2618531

Chen, F., Lu, A., Wu, H., & Li, M. (2022). Compensation and pricing strategies in cloud service SLAs: Considering participants' risk attitudes and consumer quality perception. Electronic Commerce Research and Applications, 56, 101215. https://doi.org/10.1016/j.elerap.2022.101215

Chen, H.-W., & Lim, A. (2023). A network price elasticity of demand model with product substitution. Journal of Revenue and Pricing Management, 22(4), 235–247. https://doi.org/10.1057/s41272-022-00375-w

Chen, X., Liu, Y., & Hong, G. (2023). An online learning approach to dynamic pricing and capacity sizing in service systems. Operations Research, opre.2020.612. https://doi.org/10.1287/opre.2020.612

Chiang, W.-Y. K., Chhajed, D., & Hess, J. D. (2003). Direct marketing, indirect profits: A strategic analysis of dual-channel supply-chain design. Management Science, 49(1), 1–20. https://doi.org/10.1287/mnsc.49.1.1.12749

Dan, B., Xu, G., & Liu, C. (2012). Pricing policies in a dual-channel supply chain with retail services. International Journal of Production Economics, 139(1), 312–320. https://doi.org/10.1016/j.ijpe.2012.05.014

Feng, S., Xiong, Z., Niyato, D., Wang, P., Wang, S. S., & Shen, S. X. (2022). Joint pricing and security investment in cloud security service market with user interdependency. IEEE Transactions on Services Computing, 15(3), 1461–1472. https://doi.org/10.1109/tsc.2020.2996382

Ferdous, O., Yousefi, S., & Tosarkani, B. M. (2025). A multi-disruption risk analysis system for sustainable supply chain resilience. International Journal of Disaster Risk Reduction, 116, 105136. https://doi.org/10.1016/j.ijdrr.2024.105136

Fishburn, P. C., & Odlyzko, A. M. (1999). Competitive pricing of information goods: Subscription pricing versus pay-per-use. Economic Theory, 13(2), 447–470. https://doi.org/10.1007/s001990050264

Gao, X., Zhong, W., & Mei, S. (2015). Security investment and information sharing under an alternative security breach probability function. Information Systems Frontiers, 17(2), 423–438. https://doi.org/10.1007/s10796-013-9411-3

Gera, A., & Xia, C. H. (2011). Learning curves and stochastic models for pricing and provisioning cloud computing services. Service Science, 3(1), 99–109. https://doi.org/10.1287/serv.3.1.99

Ghandour, O., El Kafhali, S., & Hanini, M. (2024). Adaptive workload management in cloud computing for service level agreements compliance and resource optimization. Computers and Electrical Engineering, 120, 109712. https://doi.org/10.1016/j.compeleceng.2024.109712

Goel, S., & Shawky, H. A. (2009). Estimating the market impact of security breach announcements on firm values. Information & Management, 46(7), 404–410. https://doi.org/10.1016/j.im.2009.06.005

Guelman, L., & Guillén, M. (2014). A causal inference approach to measure price elasticity in automobile insurance. Expert Systems with Applications, 41(2), 387–396. https://doi.org/10.1016/j.eswa.2013.07.059

Hausken, K. (2006). Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. Information Systems Frontiers, 8(5), 338–349. https://doi.org/10.1007/s10796-006-9011-6

Huang, C. D., Hu, Q., & Behara, R. S. (2008). An economic analysis of the optimal information security investment in the case of a risk-averse firm. International Journal of Production Economics, 114(2), 793–804. https://doi.org/10.1016/j.ijpe.2008.04.002

Katzan, H. (2009). Cloud software service: Concepts, technology, economics. Service Science, 1(4), 256–269. https://doi.org/10.1287/serv.1.4.256

Khajeh-Hosseini, A., Greenwood, D., Smith, J. W., & Sommerville, I. (2012). The cloud adoption toolkit: Supporting cloud adoption decisions in the enterprise. Software: Practice and Experience, 42(4), 447–465. https://doi.org/10.1002/spe.1072

Li, B., & Kumar, S. (2022). Managing software-as-a-service: Pricing and operations. Production and Operations Management, 31(6), 2588–2608. https://doi.org/10.1111/poms.13729

Li, S., Li, M., & Zhou, N. (2020). Pricing and coordination in a dual-channel supply chain with a socially responsible manufacturer. PLOS ONE, 15(7), e0236099. https://doi.org/10.1371/journal.pone.0236099

Liu, H., & Wei, S. (2022). Leveraging supply chain disruption orientation for resilience: The roles of supply chain risk management practices and analytics capability. International Journal of Physical Distribution & Logistics Management, 52(9/10), 771–790. https://doi.org/10.1108/ijpdlm-04-2021-0135

Liu, W., Yang, L., & Yu, B. (2023). Distributionally robust optimization based on kernel density estimation and mean-entropic value-at-risk. INFORMS Journal on Optimization, 5(1), 68–91. https://doi.org/10.1287/ijoo.2022.0076

Liu, Y., Meskarian, R., & Xu, H. (2017). Distributionally robust reward-risk ratio optimization with moment constraints. SIAM Journal on Optimization, 27(2), 957–985. https://doi.org/10.1137/16m106114x

Long, D. Z., & Qi, J. (2014). Distributionally robust discrete optimization with entropic value-at-risk. Operations Research Letters, 42(8), 532–538. https://doi.org/10.1016/j.orl.2014.09.004

Luna, J., Taha, A., Trapero, R., & Suri, N. (2017). Quantitative reasoning about cloud security using service level agreements. IEEE Transactions on Cloud Computing, 5(3), 457–471. https://doi.org/10.1109/tcc.2015.2469659

Luo, F., & Mehrotra, S. (2020). Distributionally robust optimization with decision dependent ambiguity sets. Optimization Letters, 14(8), 2565–2594. https://doi.org/10.1007/s11590-020-01574-3

Miller, M., & Alberini, A. (2016). Sensitivity of price elasticity of demand to aggregation, unobserved heterogeneity, price trends, and price endogeneity: Evidence from U.S. data. Energy Policy, 97, 235–249. https://doi.org/10.1016/j.enpol.2016.07.031

Nanty, S., Fiig, T., Zannier, L., & Defoin-Platel, M. (2025). Enhanced demand forecasting by combining analytical models and machine learning models. Journal of Revenue and Pricing Management, 24(5), 420–438. https://doi.org/10.1057/s41272-024-00490-w

Ojala, A. (2016). Adjusting software revenue and pricing strategies in the era of cloud computing. Journal of Systems and Software, 122, 40–51. https://doi.org/10.1016/j.jss.2016.08.070

Rahimian, H., & Mehrotra, S. (2022). Frameworks and results in distributionally robust optimization. Open Journal of Mathematical Optimization, 3, 1–85. https://doi.org/10.5802/ojmo.15

Ramachandran, M. (2016). Software security requirements management as an emerging cloud computing service. International Journal of Information Management, 36(4), 580–590. https://doi.org/10.1016/j.ijinfomgt.2016.03.008

Ranjan, A., & Jha, J. K. (2019). Pricing and coordination strategies of a dual-channel supply chain considering green quality and sales effort. Journal of Cleaner Production, 218, 409–424. https://doi.org/10.1016/j.jclepro.2019.01.297

Rose, S., & Rizopoulos, D. (2019). Machine learning for causal inference in biostatistics. Biostatistics, kxz045. https://doi.org/10.1093/biostatistics/kxz045

Ruan, J., Liu, G., Qiu, J., Liang, G., Zhao, J., He, B., & Wen, F. (2022). Time-varying price elasticity of demand estimation for demand-side smart dynamic pricing. Applied Energy, 322, 119520. https://doi.org/10.1016/j.apenergy.2022.119520

Safari, F., Safari, N., & Hasanzadeh, A. (2015). The adoption of software-as-a-service (SaaS): Ranking the determinants. Journal of Enterprise Information Management, 28(3), 400–422. https://doi.org/10.1108/jeim-02-2014-0017

Scala, B., & Lindsay, C. F. (2021). Supply chain resilience during pandemic disruption: Evidence from healthcare. Supply Chain Management: An International Journal, 26(6), 672–688. https://doi.org/10.1108/scm-09-2020-0434

Seyedan, M., & Mafakheri, F. (2020). Predictive big data analytics for supply chain demand forecasting: Methods, applications, and research opportunities. Journal of Big Data, 7(1), 53. https://doi.org/10.1186/s40537-020-00329-2

Singh, N. P., & Singh, S. (2019). Building supply chain risk resilience. Benchmarking: An International Journal, 26(7), 2318–2342. https://doi.org/10.1108/bij-10-2018-0346

Svantesson, D., & Clarke, R. (2010). Privacy and consumer risks in cloud computing. Computer Law & Security Review, 26(4), 391–397. https://doi.org/10.1016/j.clsr.2010.05.005

Wang, Q., & Ngai, E. W. T. (2022). Firm diversity and data breach risk: A longitudinal study. The Journal of Strategic Information Systems, 31(4), 101743. https://doi.org/10.1016/j.jsis.2022.101743

Wang, Z., Hayat, M. M., Ghani, N., & Shaban, K. B. (2017). Optimizing cloud-service performance: Efficient resource provisioning via optimal workload allocation. IEEE Transactions on Parallel and Distributed Systems, 28(6), 1689–1702. https://doi.org/10.1109/tpds.2016.2628370

Wibowo, B. S. (2024). When and how to adjust statistical forecasts in supply chains? Insight from causal machine learning. Journal of Business Analytics, 7(1), 25–41. https://doi.org/10.1080/2573234x.2023.2248203

Wu, C. A., Jin, C., & Liu, Q. (2024). Bundling variety, usage, or both? A multi-service analysis of pay-per-use and subscription pricing. Production and Operations Management, 33(10), 1979–1996. https://doi.org/10.1177/10591478241265480

Xin, M. (2020). The impact of customer valuation uncertainty on software licensing. MIS Quarterly, 44(2), 561–604. https://doi.org/10.25300/misq/2020/14728

Xu, X. J., Liu, Y., Ren, W. W., & Forrest, J. Y. L. (2023). Pricing decision making and coordination of a dual-channel supply chain considering online reviews and channel preferences. Managerial and Decision Economics, 44(8), 4621–4636. https://doi.org/10.1002/mde.3974

Yao, D.-Q., Wang, Z., Mukhopadhyay, S. K., & Cong, Y. (2012). Dynamic pricing strategy for subscription-based information goods. Journal of Revenue and Pricing Management, 11(2), 210–224. https://doi.org/10.1057/rpm.2010.48

Yuan, S., Das, S., Ramesh, R., & Qiao, C. (2023). Availability-aware virtual resource provisioning for infrastructure service agreements in the cloud. Information Systems Frontiers, 25(4), 1495–1512. https://doi.org/10.1007/s10796-022-10302-4

Zhang, F., & Wang, C. (2018). Dynamic pricing strategy and coordination in a dual-channel supply chain considering service value. Applied Mathematical Modelling, 54, 722–742. https://doi.org/10.1016/j.apm.2017.10.006

Zhang, Z., Nan, G., & Tan, Y. (2020). Cloud services vs. on-premises software: Competition under security risk and product customization. Information Systems Research, 31(3), 848–864. https://doi.org/10.1287/isre.2019.0919

Zhao, Y., & Liu, Q. (2023). Causal ML: Python package for causal inference machine learning. SoftwareX, 21, 101294. https://doi.org/10.1016/j.softx.2022.101294