Persuading the Supplier with Imperfect Forecast Capability

Yue Feng[1] and Tao Lu[2]

[1]College of Management and Economics, Tianjin University, Tianjin 300072

[2]School of Business, University of Connecticut, Storrs, CT 06269

September 30, 2026

Abstract

We consider a canonical supply chain in which a retailer shares its market-state forecast with an upstream supplier, who chooses production before demand realization. Because the supplier bears all overproduction costs, the retailer designs an information-sharing policy to persuade the supplier to produce more. The extant literature shows that full disclosure is optimal when state-dependent demand distributions differ only by a location shift, assuming that the retailer can perfectly predict the market state. In contrast, we focus on the same demand setting but allow for imperfect forecast capability, and show that partial or no disclosure can be optimal. When the supplier faces a high overproduction cost relative to underproduction, full disclosure is optimal at very low or high forecast capability, whereas partial disclosure can be optimal at intermediate capability. When the relative overproduction cost is low, the optimal policy progresses from no disclosure to partial disclosure and then to full disclosure as forecast capability improves. Thus, our results uncover how a retailer's forecast capability shapes its information sharing incentives. Fully sharing an imperfect forecast may be less effective than partially sharing it or not at all. The supplier can sometimes benefit from wider demand ranges associated with each state because it can induce the retailer to disclose more information.

Key words: information design; forecast sharing; supply chain management; newsvendor

1. Introduction

Advances in data infrastructure, cloud-based planning systems, and algorithmic forecasting have transformed how supply chain firms share demand information. Retailers, especially large e-commerce platforms and downstream firms close to end consumers, often possess richer demand data than their upstream suppliers. Through vendor portals, collaborative

planning systems, and algorithmic replenishment tools, these retailers can now commit to structured information-sharing protocols, allowing them to decide not only whether to share demand information, but also how much of their internal forecast should be translated into supplier-facing demand outlooks and replenishment recommendations (Liu et al. 2021, Candogan and Gurkan 2025, Chen et al. 2026).

A major challenge, however, is that retailer forecasts are themselves imperfect. Even when a retailer has superior market information, its ability to predict the true demand environment can vary substantially, for example, depending on product categories. In its collaboration with Meda, JD recognized that forecasting accuracy could be inherently constrained for some product categories. For major appliances, "inaccurate demand forecasting was a common problem facing this industry," due to low consumption frequency, difficult-to-predict repurchase timing, and large seasonal demand swings (Qi et al. 2021, p. 8). Similarly, an Amazon vendor guidance cautions that Amazon's forecasts are "usually more accurate for high-volume products in established categories and marketplaces" but "worse for long-tail selection, new products, and newly-launched countries " (Hamada 2025, sec. "How accurate are Amazon's demand forecasts?", para. 3). This observation raises the main question of this paper. How does a retailer's limited forecast capability affect its optimal information sharing strategy with an upstream supplier who must prepare production before demand is realized?

We build on the seminal supply chain setting studied in Candogan and Gurkan (2025), hereafter CG. A downstream retailer shares her information about the demand state with an upstream supplier, who must choose a production quantity before demand is realized and before the retailer places her final order. Because the supplier bears the cost of overproduction whereas the retailer benefits from higher product availability, the retailer has an incentive to persuade the supplier to produce more by designing an information-sharing policy. Our key departure from CG is that the retailer observes only an imperfect forecast of the demand state with a commonly understood accuracy. The retailer chooses how to disclose this imperfect forecast à la Bayesian persuasion (Kamenica and Gentzkow 2011).

Our main results show that imperfect forecast capability can substantially change the retailer's optimal disclosure strategy. To do so, we consider state-dependent uniform demand distributions with equal support widths, where the high-state distribution is a

rightward shift of the low-state distribution. For the same setting, assuming that the retailer can perfectly predict the demand state, CG show that full disclosure is always optimal.[1] In contrast, with imperfect forecast capability, we prove that partial or no disclosure can be optimal.

Specifically, when the supplier's critical ratio is low and he would produce within the low-state interval under the prior belief, the retailer fully discloses her forecast when forecast capability is either very high or very weak, but may partially disclose it when forecast capability is intermediate by exaggerating some low predictions as high. The intuition is that, with moderate forecast capability, fully revealing an imperfect prediction is not sufficiently persuasive to induce high enough production. Partial disclosure, by pooling some low predictions with high ones, can then be valuable because it increases the frequency of high messages, even though it reduces their credibility. When forecast capability is very weak, full disclosure is again optimal since even a favorable prediction is not sufficiently informative to induce production within the high-state interval. By contrast, when the supplier's critical ratio is high, he is already inclined to produce within the high-state interval under the prior belief. In this case, the retailer chooses no disclosure when forecast capability is low, partial disclosure when forecast capability is intermediate, and full disclosure when forecast capability is high. Additionally, when the demand supports overlap, downplaying some high predictions as low can also become optimal when forecast capability is low. Our extensive numerical studies confirm that these results are robust to nonuniform demand distributions.

Contributing to recent literature on information design in supply chains, our results advance the current understanding of a retailer's information-sharing incentives. Partial or no information sharing can arise as a consequence of imperfect forecast capability. Our findings offer the following managerial implications. Retailers should adapt their information-disclosure policies to both their own forecast capability and their suppliers' production incentives. Fully sharing a forecast with limited accuracy may be less effective than sharing it partially or not at all. Suppliers may sometimes benefit from wider demand intervals

[1] In their seminal paper, Candogan and Gurkan (2025) examine richer settings with unequal support widths, in which partial disclosure can be optimal. We intentionally focus on equal support widths to isolate the implications of imperfect forecast capability.

associated with each demand state when retailers strategically adjust disclosure, because such within-state uncertainty can induce more information sharing.

The remainder of the paper is organized as follows. §2 reviews relevant literature. §3 presents the model and §4 characterizes the retailer's optimal disclosure policy and its implications. §5 discusses more general demand settings and §6 concludes.

2. Literature

Our work is related to the large literature on demand information sharing in supply chain management. Within this literature, one stream studies better-informed upstream manufacturers sharing information with downstream retailers (Jiang et al. 2016, Jain 2022, Wang et al. 2022). More relevant to our paper is the extensive literature that studies downstream firms sharing demand information with upstream suppliers (Ha and Tong 2008). Many studies in this stream examine the impact of suppliers' strategic responses to shared information through adjustments in wholesale prices, often in competitive settings (Jain et al. 2011, Gümüş 2014, Shang et al. 2016, Huang et al. 2018).

More closely related to our paper are studies in which a downstream firm shares information about uncertain demand with an upstream supplier, who responds by making capacity or production decisions before demand realization. Cachon and Lariviere (2001) study how an informed downstream firm can credibly signal its demand forecast through contractual arrangements. Özer and Wei (2006) analyze capacity reservation and advance purchase contracts as mechanisms for enabling credible forecast sharing. Forecast sharing can occur through informal communication, giving rise to cheap-talk settings where a retailer may misreport their private information in the absence of enforceable contractual terms (Özer et al. 2011, Shamir and Shin 2016, Chu et al. 2017, Lu and Tomlin 2026). Accordingly, this literature examines whether and under what conditions credible information transmission can be sustained in equilibrium without commitment.

Our research adds to the recently growing literature on how a downstream firm can determine and commit to an ex ante information-sharing policy to influence upstream production or capacity decisions, pioneered by Candogan and Gurkan (2025). Specifically, following the Bayesian persuasion framework (Kamenica and Gentzkow 2011), Candogan and Gurkan (2025) analyze how a downstream retailer should optimally disclose its market-state forecast to an upstream supplier, who then chooses a production quantity before

demand is realized by balancing the costs of over- and underproduction. Their results shed light on how the upside and downside risks associated with state-dependent demand distributions shape the optimal information-sharing policy. Other studies examine the information-sharing policies in settings with supplier encroachment (Chen et al. 2026) and priority sourcing (Candogan et al. 2026). Inspired by and building on this literature, our work departs from Candogan and Gurkan (2025) by allowing the downstream retailer to have imperfect forecasting capability regarding the market state and asks how the retailer should optimally disclose its imperfect forecast. As detailed later, our analysis shows that the retailer may partially disclose or withhold its forecast when state-dependent demand follows uniform distributions that differ only by a location shift – a setting in which Candogan and Gurkan (2025) show that full disclosure is always optimal when the retailer can perfectly forecast the market state.

More broadly, information design problems have been studied in other contexts, such as the disclosure of inventory availability (Drakopoulos et al. 2021), the provision of product quality information (Boleslavsky et al. 2017, Boyaci et al. 2024), platform operations (Küçükgül et al. 2022, Bimpikis et al. 2024), agricultural operations (Farahani et al. 2025), and public policy settings (Alizamir et al. 2020, De Véricourt et al. 2021, Xu et al. 2026); see Candogan (2020) for a survey of the literature in operations management. To our knowledge, only a few papers have considered Bayesian persuasion problems in which the sender has imperfect capability to reveal the true state. Lu and Tomlin (2025) study the strategic disclosure of a machine-generated demand forecast to incentivize a human operations manager to exert additional forecasting effort that complements the machine forecast. Yang et al. (2026) examine the design of a disinformation detector based on an imperfect classifier when senders strategically choose whether to generate disinformation. In contrast, the incentive misalignment in our model arises from the supplier bearing the overstocking risk while the retailer benefits from greater product availability, giving rise to distinct economic tradeoffs that shape the optimal disclosure policy.

Our focus on forecast capability also connects to the literature examining the operational implications of demand-forecast accuracy. Iyer et al. (2007) study how an imperfect demand signal, jointly observed by a manufacturer and a retailer, affects inventory, pricing, and contracting decisions in a distribution channel. Taylor and Xiao (2010) examine the impact

of a retailer's forecast accuracy when an upstream manufacturer can design a menu of screening contracts to elicit the retailer's private forecast. Chen and Xiao (2012) analyze the impact of a reseller's forecast accuracy in a three-tier supply chain consisting of an upstream manufacturer, the reseller, and a salesperson. Our Bayesian persuasion setting differs substantially from these studies, and our results provide new insights into how a retailer's forecast accuracy shapes its incentives for disclosing its private forecast.

3. Model

3.1. Supply Chain Primitives

We consider a canonical supply chain setting, as in CG. A downstream retailer (she) procures a product from an upstream supplier (he) to satisfy uncertain market demand D. Due to long lead times, the supplier must choose a production quantity Q before the retailer observes realized demand and places her final order. The supplier can therefore fill the retailer's order only up to its available production quantity, delivering $\min(D, Q)$ units.[2] The retailer sells the product at retail price p, and the supplier incurs a production cost c per unit. The retailer pays the supplier a given wholesale price w for each unit delivered. We assume that $p > w > c$, and that any leftover unit has zero salvage value.

As in CG, demand follows one of two probability distributions, depending on the market state $\omega \in \{H, L\}$. Specifically, we consider $D_\omega = \theta_\omega + \epsilon$ where $\theta_H > \theta_L$. Thus, $\omega = H$ and $\omega = L$ represent a high and low market state, respectively. The term ϵ captures the uncertainty conditional on each market state. For notational brevity, we assume that the two demand states occur with equal prior probabilities; that is, $\Pr(\omega = H) = \Pr(\omega = L) = \frac{1}{2}$.[3] In line with CG, we consider uniformly distributed demand under each market state by assuming that ϵ follows a uniform distribution $U(-\sigma, \sigma)$, with $0 < \sigma \leq \theta_L$ to ensure non-negative demand. Under this specification, demand differs across market states only in its mean, while the support size (and hence the variance) is the same across states. We adopt this specification because CG has explored how state-dependent support sizes shape the retailer's information-disclosure policy, assuming that the retailer can perfectly forecast the market state. In the present paper, we intentionally hold the support size identical across

[2] Alternatively, one may interpret Q as the production capacity that the supplier must build in advance.

[3] The equal-prior assumption has been widely adopted in the demand information-sharing literature; see, e.g., Iyer et al. (2007) and Jiang et al. (2016). Allowing unequal prior probabilities does not affect our main results, but only complicates the mathematical expressions.

states, which enables us to isolate the effect of the retailer's imperfect forecast capability on the optimal disclosure policy.

We define $\Delta := \theta_H - \theta_L$ and focus on the case $\Delta \geq 2\sigma$ as our base model, where the demand supports under the low and high market states do not overlap. We analyze the case $\Delta < 2\sigma$ in §5.1 and further examine the case of nonuniform demand in §5.2.

3.2. Information Structure

Because the retailer is closer to the end market, she possesses superior information about the market state. The retailer can transform this internal information into supplier-facing reports (such as demand outlooks or replenishment recommendations), and share them through a vendor data-sharing platform (Qi et al. 2021). As in CG, we model the retailer as choosing an information disclosure policy that maps her internal forecast into a message sent to the supplier. Importantly, we depart from CG by assuming that the retailer has *imperfect* forecast capability, so that her internal forecast may or may not coincide with the actual market state.

Specifically, the retailer receives a signal $X \in \{0, 1\}$ as her prediction of the market state ω, where $X = 1$ and $X = 0$ indicate that the signal predicts the market state to be high and low, respectively. Let $\alpha \in (\frac{1}{2}, 1]$ denote the accuracy of X, defined as

$$\alpha = \Pr(X = 1 \mid \omega = H) = \Pr(X = 0 \mid \omega = L),$$

that is, the probability that X correctly predicts the market state conditional on each realization of ω.[4] A larger α signifies stronger forecast capability of the retailer. As $\alpha \to \frac{1}{2}$, the signal X becomes uninformative. When $\alpha = 1$, X perfectly reveals the market state. As discussed in the introduction, it is widely recognized that demand for certain product categories, such as newly launched products or products with low purchase frequency, is harder to predict. We therefore assume that the retailer's forecast capability, α, is common knowledge.[5] The signal X is independent of the noise term ϵ conditional on the market state ω.

[4] Together with the equal-prior assumption, this symmetric accuracy structure implies that X is an unbiased estimator of ω, that is, $\Pr(X = 1) = \Pr(\omega = H) = \frac{1}{2}$ and $\Pr(X = 0) = \Pr(\omega = L) = \frac{1}{2}$.

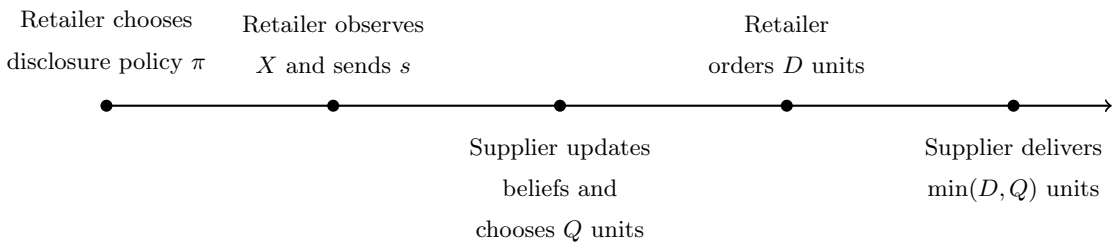
[5] In the supply chain literature, it is rather common to assume that the retailer has a private forecast, whereas the accuracy of the forecast is common knowledge (e.g., Taylor and Xiao 2010, Chen and Xiao 2012).

The retailer chooses a disclosure policy that maps each realization of X into a probability distribution over messages $s \in M$, where $M = \{h, l\}$ is a message space. Message $s = h$ ($s = l$) denotes a high (low) message and can alternatively be interpreted as a recommendation for a high (low) production quantity.[6] In other words, the retailer determines a collection of conditional probabilities $\{\pi(s \mid X)\}_{X \in \{0,1\}}$, where $\pi(s \mid X)$ denotes the probability of sending message s conditional on the realized forecast X. We define the shorthand notation $\pi_h = \pi(h \mid 1)$ and $\pi_l = \pi(l \mid 0)$ to represent the probabilities that the message matches the realized forecast X. A policy is full disclosure if $\pi_h = \pi_l = 1$, so that the message s fully reveals X. It is no disclosure if the distribution of messages is independent of X. For expositional convenience, we represent no disclosure by $\pi_h = 1$ and $\pi_l = 0$, under which the retailer always sends the high message regardless of X.[7] A policy is said to be partial disclosure if it is neither full disclosure nor no disclosure. Under a partial disclosure policy, the retailer may possibly send a message inconsistent with the realized forecast with $\pi_h < 1$ or $\pi_l < 1$.

The sequence of events, as illustrated by Figure 1, is as follows:

1. The retailer determines an information disclosure policy π.
2. The retailer observes X, and sends a message s, following π.
3. The supplier updates his belief, and chooses the production quantity Q.
4. Demand D is realized, and the retailer places an order of D units.
5. The supplier delivers $\min(D, Q)$ units, and both parties receive their payoffs.

Figure 1 Timing of Events



[6] Restricting attention to a binary message space is without loss of optimality, as a binary message space is sufficient to attain the optimum in our setting.

[7] More generally, no disclosure requires $\pi_h + \pi_l = 1$, so that the distribution of messages is independent of X. For example, $\pi_h = 0$ and $\pi_l = 1$ also represents no disclosure, under which the retailer always sends the low message.

As in CG, the retailer announces and commits to a disclosure policy π at the outset. This commitment assumption is reasonable in settings where supply chain members engage in long-term collaboration and repeated interactions (Best and Quigley 2024). It is also consistent with modern supply chain relationships supported by digital and cloud-based systems, which allow retailers to implement pre-specified information-sharing protocols based on algorithmic forecasts and analytics (Qi et al. 2021, Hamada 2025). We refer interested readers to CG and the recent literature in supply chain information design (e.g., Chen et al. 2026) for a more detailed justification of the commitment assumption.

In typical information-design settings, the sender chooses among all possible information structures that directly garble the payoff-relevant state. In our model, however, the retailer herself only has access to a noisy signal X of the market state ω. In this regard, the retailer's disclosure policy is constrained by the informativeness of X: she cannot reveal the market state beyond X, but can choose how to convey the information contained in X to disclose to the supplier.

Let $\mu(s) = \Pr(\omega = H \mid s)$ denote the supplier's posterior belief that the market state is high after receiving message s. For simplicity, we often write this posterior as μ, with the understanding that it depends on the realization of s. Since the message s is generated from the retailer's forecast X, the supplier forms his posterior belief about ω through his posterior belief about X. By the law of total probability,

$$\begin{aligned} \mu &= \Pr(\omega = H \mid X = 1)\ \Pr(X = 1 \mid s) + \Pr(\omega = H \mid X = 0)\ \Pr(X = 0 \mid s) \\ &= \alpha\tau + (1-\alpha)(1-\tau), \end{aligned} \tag{1}$$

where τ represents the supplier's posterior belief regarding $X = 1$ and by Bayes' rule,

$$\tau = \Pr(X = 1 \mid s) = \frac{\pi(s|1)\Pr(X = 1)}{\pi(s|1)\Pr(X = 1) + \pi(s|0)\Pr(X = 0)}. \tag{2}$$

By choosing the disclosure policy π, the retailer influences the posterior belief τ induced by each message, subject to Bayes plausibility across messages. Because τ take values in $[0, 1]$, the set of *inducible* market-state posteriors is restricted to a subset of $[0, 1]$:

$$\mu = \alpha\tau + (1-\alpha)(1-\tau) \in [1-\alpha, \alpha].$$

For instance, even when the retailer fully discloses her forecast X and the forecast is high, $X = 1$, the supplier recognizes that the forecast is imperfect and therefore forms

the posterior belief $\mu = \alpha$, rather than $\mu = 1$. This echoes practical examples in which, for product categories commonly recognized as having low demand predictability, such as infrequently purchased major appliances or newly launched products, suppliers may remain skeptical of a retailer's forecast even when the retailer fully shares it.

REMARK 1. While we describe the retailer's problem as conveying or disclosing her forecast X, our model can alternatively be interpreted as a problem of choosing the forecast precision to implement. Specifically, X can be viewed as the most informative prediction generated by the best available forecasting engine, whose true positive and true negative rates are both equal to α. By choosing a disclosure policy π, the retailer can implement any information structure (i.e., a set of true positive and true negative rates) that is weakly less informative than X in the Blackwell sense, because the message s is a garbling of X (Blackwell 1953). From a practical implementation perspective, partial disclosure of X can thus be interpreted as the retailer choosing to implement a forecasting engine that is less informative than the best available one.

3.3. Profit Functions

Because the retailer's disclosure policy π and the realized message s influence the supplier's production decision only through its posterior belief μ, we formulate the supplier's problem as a function of μ. Given any posterior belief μ regarding the market state, the supplier chooses a production quantity Q to maximize his expected profit

$$\Pi_S(Q) = w\mathbb{E}\big[\mu \min(\theta_H + \epsilon, Q) + (1-\mu)\min(\theta_L + \epsilon, Q)\big] - cQ, \tag{3}$$

where the expectation is taken over ϵ. Let Q^*_μ denote the optimal solution to problem (3) given posterior belief μ.

Anticipating the supplier's best response Q^*_μ, the retailer chooses π to maximize its expected profit. Let

$$\Pi_R(\mu) = (p-w)\mathbb{E}\left[\mu \min(\theta_H + \epsilon, Q^*_\mu) + (1-\mu)\min(\theta_L + \epsilon, Q^*_\mu)\right]$$

denote the retailer's expected payoff when the supplier holds posterior belief μ and optimally chooses the production quantity. As reflected in the profit functions above, the supplier bears the overproduction risk and may therefore be reluctant to produce a large

quantity, whereas the retailer always benefits from a higher production quantity. This creates an incentive misalignment between the supplier and the retailer. To persuade the supplier to produce a sufficiently large quantity, the retailer chooses a disclosure policy π at the beginning of the game that solves the following problem:

$$\max_{\pi} \mathbb{E}_X \mathbb{E}_{s\sim\pi(\cdot|X)} \left[\Pi_R(\mu(s))\right], \tag{4}$$

where $\mu(s)$ is determined by Bayes' rule according to (1) and (2).

As shown by Kamenica and Gentzkow (2011), the above problem can equivalently be formulated as choosing a distribution over posterior beliefs μ subject to Bayes plausibility. However, as discussed above, when $\alpha < 1$, the market-state posterior μ is restricted to a strict subset of $[0, 1]$. As a result, the standard concavification approach with respect to μ does not apply directly. In the analysis, we instead optimize over τ, the supplier's posterior belief about the retailer's forecast X, since τ spans the full range $[0, 1]$. With this technical point in mind, the subsequent discussion continues to focus on μ, since the market-state posterior provides the more intuitive basis for interpretation.

4. Results

4.1. Optimal Production Quantity and Retailer Profit Given Posterior μ

We begin by characterizing the supplier's optimal production quantity for any given posterior belief μ about the market state being high, with the understanding that μ is restricted to $[1-\alpha, \alpha] \subseteq [0, 1]$. We define $r = \frac{w-c}{w}$ as the supplier's critical ratio. As is standard in newsvendor problems, a larger r means that the supplier faces a higher underage cost, $w - c$, relative to the overage cost, c, and therefore tends to produce a larger quantity.

LEMMA 1. *For any belief μ, the supplier's optimal production quantity Q^*_μ is given by*

$$Q^*_\mu = \begin{cases} \theta_L + \sigma\left(\frac{2r}{1-\mu} - 1\right) & \text{if } \mu < 1 - r, \\ \theta_H + \sigma\left(1 - \frac{2(1-r)}{\mu}\right) & \text{if } \mu \geq 1 - r. \end{cases} \tag{5}$$

Moreover, Q^*_μ is increasing in μ, convex in μ when $\mu < 1 - r$, and concave in μ when $\mu \geq 1 - r$.[8]

[8] When multiple production quantities are optimal for the supplier, we break ties in favor of the larger quantity.

The supplier's optimal production quantity Q_μ^* falls into one of two cases, depending on the critical ratio r and the posterior belief μ. When $\mu < 1-r$, the supplier believes that the high-demand state is unlikely and therefore chooses a production quantity within the low-state demand interval, i.e., $Q_\mu^* \in [\theta_L - \sigma, \theta_L + \sigma]$. When $\mu \geq 1-r$, the supplier is sufficiently optimistic about the high-demand state and chooses a production quantity within the high-state demand interval, i.e., $Q_\mu^* \in [\theta_H - \sigma, \theta_H + \sigma]$. Moreover, the curvature properties of Q_μ^* are useful for later discussions. The supplier clearly chooses a larger quantity as his posterior belief μ increases. The second-order effect, however, is more nuanced. When $\mu < 1-r$ and the supplier produces within the low-state demand interval, Q_μ^* is convex in μ. This suggests that an increase in the high-state posterior has an increasing marginal effect in inducing the supplier to choose a larger production quantity when he is still targeting low-state demand. In contrast, when $\mu \geq 1-r$ and the supplier produces within the high-state demand interval, Q_μ^* is concave in μ. This suggests that an increase in the posterior has a diminishing marginal effect on the supplier's production quantity once he is already targeting high-state demand.

Invoking the expressions of Q_μ^*, we can derive the retailer's profit as a function of the supplier's posterior μ.

LEMMA 2. *The retailer's expected profit as a function of the belief μ is*

$$\Pi_R(\mu) = \begin{cases} (p-w)\left[\theta_L + \dfrac{\sigma\big(\mu - (1-r)^2\big)}{1-\mu}\right] & \text{if } \mu < 1-r, \\ (p-w)\left[\theta_L + \Delta\mu - \dfrac{\sigma(1-r)^2}{\mu}\right] & \text{if } \mu \geq 1-r. \end{cases}$$

Moreover, $\Pi_R(\mu)$ is increasing in μ, convex in μ when $\mu < 1-r$ and concave in μ when $\mu \geq 1-r$.

Lemma 2 confirms that the retailer's profit $\Pi_R(\mu)$ inherits the same monotonicity and convex-concave structure as the supplier's optimal production quantity Q_μ^*, since the retailer always benefits from a larger supplier production quantity.

The curvature of Q_μ^* and $\Pi_R(\mu)$ has implications for information disclosure. Because the induced distribution of posteriors must satisfy Bayes plausibility, an information-disclosure policy affects the retailer's expected payoff by changing the dispersion of μ in a mean-preserving manner. Thus, Jensen's inequality suggests a local benefit from generating

dispersion in posterior beliefs when $\Pi_R(\mu)$ is convex, and a local benefit from pooling posterior beliefs when $\Pi_R(\mu)$ is concave. Globally, the optimal disclosure policy depends on the relative importance of these convex and concave regions.

4.2. Optimal Information Disclosure Policy

4.2.1. Perfect Forecast Capability. When the retailer can perfectly predict the market state, i.e., $\alpha = 1$, our model reduces to the disjoint-support case studied in Theorem 3 of CG. In particular, CG show that full disclosure is optimal when the demand supports under the two market states have equal size. We restate their result below as a benchmark.[9]

PROPOSITION 1. *(Candogan and Gurkan 2025, Theorem 3) When $\alpha = 1$, full disclosure is optimal; that is, $\pi_h^* = \pi_l^* = 1$.*

Because the retailer does not bear the supplier's overproduction cost, she always prefers the supplier to produce more. It is intuitive that the retailer fully reveals high-state predictions by setting $\pi_h = 1$. Yet, it seems that the retailer might benefit from exaggerating low-state predictions as high with positive probability, by setting $\pi_l < 1$, or equivalently $\pi(h \mid X = 0) > 0$. However, doing so reduces the credibility of the high-state message $s = h$ and may lower the supplier's production response, because the supplier forms his posterior belief with consideration of the disclosure policy π. There is a potential tradeoff between maintaining the credibility of high messages and sending high messages more frequently. The result of CG implies that, when the high- and low-state demand distributions have equal support widths and the retailer has perfect forecast capability, the benefit of maintaining credibility through full disclosure always dominates the gain from increasing the frequency of sending high messages. This point is illustrated by Example 1.

EXAMPLE 1. Table 1 presents a numerical example in which the retailer has perfect forecast capability, $\alpha = 1$, and the supplier has a critical ratio $r = 0.4$, with $\theta_H = 6$, $\theta_L = 1$, and $\sigma = 1$. Under full disclosure, messages l and h are sent with equal probabilities, i.e., $\Pr(l) = \Pr(h) = 1/2$. Since X is perfect, the resulting posteriors are $\mu(l) = 0$ and $\mu(h) = 1$, leading to $Q^*_{\mu(l)} = 0.80$ and $Q^*_{\mu(h)} = 5.80$, respectively.

[9] Candogan and Gurkan (2025) also consider cases in which the demand supports differ across market states. In those cases, partial or no disclosure can be optimal when the downside support is sufficiently large relative to the upside support. We abstract from this support asymmetry to isolate the implications of imperfect forecast capability.

Now consider the partial-disclosure policy $\pi_h = 1$ and $\pi_l = 2/3$, under which the retailer always sends a high message following a high forecast but exaggerates a low forecast by sending a high message with probability $1 - \pi_l = 1/3$. This increases the probability of sending a high message from $\Pr(h) = 1/2$ to $\Pr(h) = 2/3$. Meanwhile, it also weakens the credibility of the high message. By Bayes' rule, the supplier's posterior following a high message falls from $\mu(h) = 1$ to $\mu(h) = \frac{(1/2)\times 1}{(1/2)\times 1 + (1/2)\times(1/3)} = \frac{3}{4}$, resulting in a lower production quantity, $Q^*_{\mu(h)} = 5.40$. Comparing the retailer's expected payoffs, $\frac{1}{2}\Pi_R(0) + \frac{1}{2}\Pi_R(1) = 628 > \frac{1}{3}\Pi_R(0) + \frac{2}{3}\Pi_R\left(\frac{3}{4}\right) = 612$, shows that exaggerating low messages to increase the frequency of high messages does not outperform full disclosure which preserves the maximum credibility of a high message.

Table 1 Full versus partial disclosure under perfect forecast capability

α	Disclosure Policy	π_l	π_h	$\Pr(l)$	$\Pr(h)$	$\mu(l)$	$\mu(h)$	$Q^*_{\mu(l)}$	$Q^*_{\mu(h)}$	$\mathbb{E}[\Pi_R(\mu)]$
1.00	Full*	**1**	**1**	**1/2**	**1/2**	**0.00**	**1.00**	**0.80**	**5.80**	**628.00**
	Partial	2/3	1	1/3	2/3	0.00	0.75	0.80	5.40	612.00

Note: Parameter values are $\theta_H = 6$, $\theta_L = 1$, $\sigma = 1$, $p = 1000$, $w = 800$, and $c = 480$, so that $r = 0.4$.

However, the above comparison is based on $\alpha = 1$. As we show below, the partial disclosure policy $\pi_h = 1$ and $\pi_l = \frac{2}{3}$ is in fact optimal when $\alpha = 0.7$, with all else unchanged; that is, the tradeoff between message frequency and credibility can reverse when the retailer's forecast capability is imperfect.

4.2.2. Imperfect Forecast Capability. We now allow $\alpha < 1$, so that the retailer's forecast X can be noisy. As a result, even if the retailer wishes to fully reveal the market state ω, she lacks the capability to do so perfectly. We characterize the retailer's optimal disclosure policy in three cases, depending on the supplier's production response under the inducible posterior beliefs: whether the supplier can ever be induced to produce within the high-state demand interval, and whether he would do so even under the prior belief. To streamline notation, we define several thresholds that are useful for characterizing the solution structure: $\alpha^{\mathrm{C}} = \frac{\Delta + 2r\sigma + \sqrt{\Delta^2 - 4r(1-r)\sigma(\Delta+2\sigma)}}{2(\Delta+\sigma)}$, $\alpha^{\mathrm{N}} = \frac{\Delta + \sigma + \sqrt{(\Delta+\sigma)^2 - 4[\Delta + 4\sigma(1-r)^2]\sigma r(2-r)}}{2[\Delta + 4\sigma(1-r)^2]}$, and α^{F} is the largest root in $[1/2, 1)$ of $\Delta\alpha^3 - (\Delta+\sigma)\alpha^2 + \sigma\left[1 + 2(1-r)^2\right]\alpha - \sigma(1-r)^2 = 0$.

Proposition 2. *When $\alpha \in (1/2, 1)$, the retailer's optimal information policy is characterized as follows:*

- *Case* I*: If $r < 1-\alpha$, full disclosure is optimal; that is, $\pi_h^* = \pi_l^* = 1$.*
- *Case* II*: If $1-\alpha \le r < \frac{1}{2}$, the optimal policy is characterized as follows.*

  (a) If $\alpha \ge \alpha^{\mathrm{F}}$, full disclosure is optimal; that is, $\pi_h^* = \pi_l^* = 1$.

  (b) If $\alpha^{\mathrm{C}} < \alpha < \alpha^{\mathrm{F}}$, a balanced partial disclosure policy is optimal with $\pi_h^* = 1$ and $\pi_l^* = 2 - \frac{1}{\tau^o}$ where

$$\tau^o = \left(1 - \frac{\alpha}{1-\alpha} + \frac{\alpha}{1-\alpha} \cdot \frac{(2\alpha-1)(1-r)\sigma}{\sqrt{\alpha\sigma\left[\sigma\left((1-\alpha)^2 + (2\alpha-1)(1-r)^2\right) - (1-\alpha)^2\alpha\Delta\right]}}\right)^{-1}. \tag{6}$$

  (c) Otherwise, a cutoff partial disclosure policy is optimal with $\pi_h^* = 1$ and $\pi_l^* = 2 - \frac{1}{\tau^\dagger}$ where $\tau^\dagger = \frac{\alpha - r}{2\alpha - 1}$.
- *Case* III*: If $r \ge \frac{1}{2}$, the optimal policy is characterized as follows.*

  (a) If $\alpha \ge \alpha^{\mathrm{F}}$, full disclosure is optimal with $\pi_h^* = \pi_l^* = 1$.

  (b) If $\alpha^{\mathrm{N}} < \alpha < \alpha^{\mathrm{F}}$, the balanced partial disclosure policy is optimal with $\pi_h^* = 1$ and $\pi_l^* = 2 - \frac{1}{\tau^o}$ where τ^o is defined in (6)*.*

  (c) If $\alpha \le \alpha^{\mathrm{N}}$, no disclosure is optimal with $\pi_h^* = 1$ and $\pi_l^* = 0$.[10]

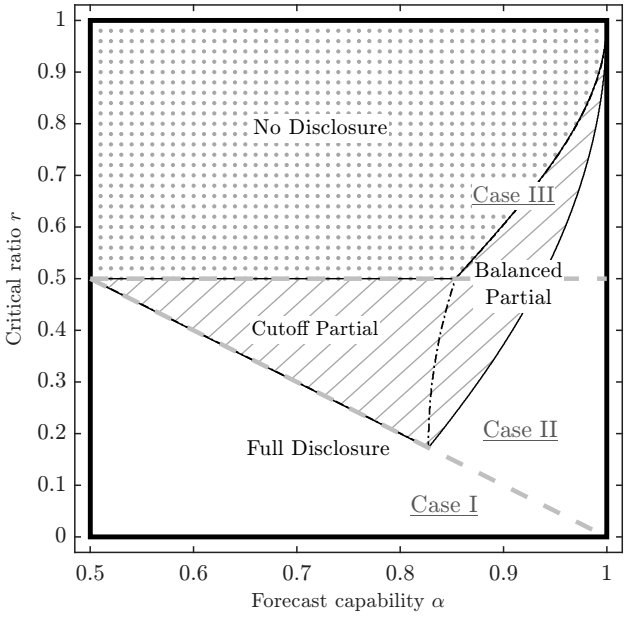
The structure of optimal disclosure policies is illustrated by Figure 2, where the two gray dashed lines partition the r-α space into Cases I, II and III, as characterized in Proposition 2.

Case I. When $r < 1-\alpha$, both the supplier's critical ratio and the retailer's forecast capability are sufficiently low that the supplier always produces within the low-state demand interval for all inducible posteriors. Indeed, by Lemma 1, the supplier produces for the high state only when $\mu \ge 1-r$. However, due to the retailer's limited forecast capability, the highest posterior she can induce is $\mu = \alpha$, attained when she fully discloses the high signal $X = 1$. Since $\alpha < 1-r$, the retailer cannot convince the supplier to produce within the high-state demand interval. As shown earlier, the retailer's payoff is convex in μ throughout the low-state region. Therefore, full disclosure maximizes the retailer's expected payoff. Case I corresponds to the lower-left triangular region in Figure 2.

Case II. When $1-\alpha \le r < \frac{1}{2}$, the supplier can be induced to produce within the high-state demand interval if the retailer's message generates a posterior belief $\mu \ge 1-r$, which

[10] As mentioned, we represent no disclosure by setting $\pi_h = 1$ and $\pi_l = 0$, so that the retailer always sends message h and the supplier therefore learns nothing from the message and relies on the prior. Other representations, such as $\pi_h = 0$ and $\pi_l = 1$, can also implement no disclosure.

Figure 2 Structure of the Optimal Policy



Note. $\theta_H = 6$, $\theta_L = 1$, $\sigma = 1$.

is feasible because $\alpha \geq 1 - r$. However, if no additional information is provided beyond the prior, the supplier still produces within the low-state demand interval because $r < \frac{1}{2}$. Now the optimal disclosure policy falls into three subcases depending on the forecast capability α. In subcase (a), α is sufficiently close to one, i.e., $\alpha \geq \alpha^F$, and full disclosure is optimal, as in CG. In this case, preserving the credibility of the high-state message is more valuable than increasing its frequency by pooling it with low-state predictions.

The key reversal occurs when α falls below α^{F} as in subcases (b) and (c). Because the retailer's forecast is less accurate, fully revealing a high-state prediction no longer induces a sufficiently attractive production response. The retailer then benefits more from increasing the frequency of high-state messages, even though doing so reduces their credibility.

Specifically, in subcase (b), the retailer always sends a high message following a high-state prediction, i.e., $\pi_h^* = 1$, but exaggerates a low-state prediction by sending a high message with probability $1 - \pi_l^* = \frac{1}{\tau^o} - 1 > 0$. Here, τ^o denotes the supplier's posterior belief that the retailer's prediction is $X = 1$ conditional on a high message under the optimal disclosure policy. Compared with full disclosure, under which a high message is generated with probability $\Pr(h) = \frac{1}{2}$ and induces a posterior $\mu(h) = \alpha$, the optimal partial disclosure policy trades off credibility for frequency: A realized high message becomes less credible and

thus induces a lower production quantity, because upon receiving $s = h$, the supplier forms a posterior $\mu(h) = \alpha\tau^o + (1-\alpha)(1-\tau^o) < \alpha$. On the other hand, the probability of sending a high message increases to $\Pr(h) = \frac{1}{2\tau^o} > \frac{1}{2}$. Intuitively, this tradeoff becomes beneficial because, with imperfect forecast capability α, even full disclosure induces a posterior belief of at most α. Compared to the case of perfect forecast capability, exaggerating low-state predictions causes a smaller loss in credibility, and hence a smaller reduction in the resulting production quantity, which can be outweighed by the benefit of increasing the frequency of high messages. The value of τ^o balances this credibility-frequency tradeoff in the interior of the high-production region. We therefore refer to the policy in subcase (b) as the *balanced partial disclosure* policy.

In subcase (c), α becomes even smaller, making it more attractive to increase the frequency of high messages by exaggerating low-state predictions more often. The optimal policy in fact maximizes the probability of sending a high message while keeping the induced posterior just high enough for the supplier to produce within the high-state demand interval; that is, $\tau^\dagger$ is chosen such that $\mu(h) = 1 - r$ where $\tau^\dagger$ denotes the supplier's posterior belief that $X = 1$ conditional on a high message. We therefore refer to the policy in subcase (c) as the *cutoff partial disclosure* policy.

EXAMPLE 2. Let us now revisit Example 1. Assume that all the parameters remain unchanged except that the retailer has imperfect forecast capability $\alpha = 0.7$.

Table 2 Full versus partial disclosure under imperfect forecast capability ($r = 0.4$)

α	Disclosure Policy	π_l	π_h	$\Pr(l)$	$\Pr(h)$	$\mu(l)$	$\mu(h)$	$Q^*_{\mu(l)}$	$Q^*_{\mu(h)}$	$\mathbb{E}[\Pi_R(\mu)]$
0.70	Full	1	1	1/2	1/2	0.30	0.70	1.14	5.29	490.00
	Partial*	**2/3**	**1**	**1/3**	**2/3**	**0.30**	**0.60**	**1.14**	**5.00**	**514.29**

Note: Parameter values are $\theta_H = 6$, $\theta_L = 1$, $\sigma = 1$, $p = 1000$, $w = 800$, and $c = 480$, so that $r = 0.4$.

Substituting the parameter values into Proposition 2, one can show that the cutoff partial-disclosure policy described in Case II(c) is optimal, with $\pi_h^* = 1$ and $\pi_l^* = 2/3$. Thus, the policy that is suboptimal under perfect forecast capability, $\alpha = 1$, as shown in Example 1, becomes optimal when $\alpha = 0.7$. A comparison of Tables 1 and 2 shows that, by exaggerating low-state predictions with probability 1/3, the partial disclosure policy increases the probability of a high message from 1/2 to 2/3 in both cases. However, the loss of credibility

is smaller when $\alpha = 0.7$: the posterior $\mu(h)$ falls only from 0.70 to 0.60, compared with a decline from 1.00 to 0.75 when $\alpha = 1$.[11] Consequently, the reduction in the supplier's production response to a high message is also smaller, making it more attractive for the retailer to increase the frequency of high messages at the expense of some credibility.

Case III. When $r \geq \frac{1}{2}$, the critical ratio is sufficiently high that the supplier already produces within the high-state demand interval under the prior belief. Similar to Case II, when α is sufficiently high, the benefit of preserving message credibility dominates, and full disclosure is optimal. For intermediate values of α, the retailer balances credibility and frequency through the balanced partial-disclosure policy. However, unlike Case II, when α is sufficiently low, no disclosure is optimal rather than full disclosure. Because the prior already induces high-state production, revealing an inaccurate prediction provides only a limited production increase following a high prediction but causes a production reduction following a low prediction. This intuition is captured by the concavity of the retailer's payoff within the high-production region, making no disclosure optimal.

EXAMPLE 3. Table 3 presents a numerical example with supplier critical ratio $r = 0.7$. Under no disclosure, the supplier effectively ignores the message and produces 5.80 based on the prior. The optimal policy switches sharply from full disclosure to no disclosure as the retailer's forecast capability α decreases only slightly, from 1 to 0.9.

Table 3 Full versus no disclosure under perfect and imperfect forecast capabilities ($r = 0.7$)

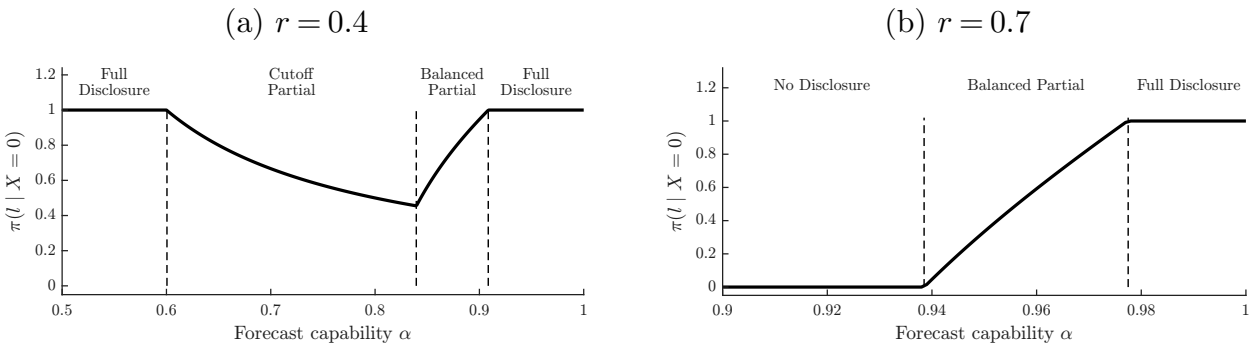
α	Policy	π_l	π_h	$\Pr(l)$	$\Pr(h)$	$\mu(l)$	$\mu(h)$	$Q^*_{\mu(l)}$	$Q^*_{\mu(h)}$	$\mathbb{E}[\Pi_R(\mu)]$
1.00	Full*	**1**	**1**	**0.5**	**0.5**	**0**	**1**	**1.40**	**6.40**	**682.00**
	No Disclosure	0	1	0	1	0.5	0.5	5.80	5.80	664.00
0.90	Full	1	1	0.5	0.5	0.1	0.9	1.56	6.33	641.11
	No Disclosure*	**0**	**1**	**0**	**1**	**0.5**	**0.5**	**5.80**	**5.80**	**664.00**

Note: Parameter values are $\theta_H = 6$, $\theta_L = 1$, $\sigma = 1$, $p = 1000$, $w = 800$, and $c = 240$.

Overall, our results suggest that when the retailer has imperfect forecast capability, she should tailor the information-sharing policy to her forecast accuracy as well as the supplier's incentive to produce. Specifically, with a given $\alpha < 1$, the retailer should fully disclose her

[11] Generally speaking, the market-state posterior $\mu = \alpha\tau + (1-\alpha)(1-\tau)$ becomes less sensitive to the retailer-forecast posterior $\tau = \Pr(X = 1 \mid s)$ as α decreases, because $\partial\mu/\partial\tau = 2\alpha - 1$. Thus, for a given probability of exaggerating low-state predictions, the resulting loss in the market-state posterior is smaller, while the increase in the frequency of high messages remains the same.

Figure 3 Optimal policy π_l^* as a function of α.



Note. $\theta_H = 6$, $\theta_L = 1$, $\sigma = 1$.

forecast if the supplier's critical ratio is low, partially disclose it if the supplier's critical ratio is moderate, and disclose no information if the supplier's critical ratio is sufficiently high. It is worth noting that when the retailer's forecast capability is perfect and the state-dependent demand distributions differ in support widths, CG identify a broad range of settings in which the optimal disclosure policy switches from partial disclosure to no disclosure and then to full disclosure as the supplier's critical ratio increases – a pattern substantially different from ours. This difference suggests that partial disclosure driven by imperfect forecast capability operates through a distinct mechanism from that driven by unequal demand support widths.

4.3. Effect of Forecast Capability α

From the structure of the optimal disclosure policies, it follows that, with all else fixed, as the retailer's forecast capability α improves, the retailer may disclose either more or less information about her inherent prediction X. Since the retailer always fully reveals high-state predictions, i.e., $\pi_h^* = 1$, the key measure of disclosure is π_l^*, the probability that the retailer truthfully reveals a low-state prediction. A higher π_l^* means that the retailer discloses more information about her inherent prediction. The proposition below characterizes when π_l^* is monotone or non-monotone in α. Throughout the paper, unless otherwise specified, we use "increasing" and "decreasing" in the weak sense.

PROPOSITION 3. *(i) Given $r < \frac{1}{2}$, π_l^* is decreasing in α if $\alpha \le \alpha^C$, and increasing in α if $\alpha > \alpha^C$. (ii) Given $r \ge \frac{1}{2}$, π_l^* is increasing in α.*

The optimal π_l^* can be non-monotone in α when $r < 1/2$. In this case, the supplier produces within the low-state demand interval under the prior belief, while the retailer

attempts to persuade him to produce within the high-state demand interval. In contrast, π_l^* is increasing in α when $r \geq 1/2$, in which case the supplier already produces within the high-state demand interval under the prior belief. Figure 3 numerically illustrates the two cases with $r = 0.4$ and $r = 0.7$, respectively.

First, we consider the case of $r = 0.4$ in Figure 3(a). For any $\alpha < 1 - r = 0.6$, the retailer cannot induce the supplier to produce within the high-state demand interval, even by fully disclosing a high-state prediction. This corresponds to Case I of the optimal policy characterization, where the retailer chooses full disclosure, $\pi_l^* = 1$. Under this policy, the supplier always produces within the low-state demand interval, although the exact production quantity Q^* varies with the realized prediction.

As α increases beyond 0.6, the setting enters Case II of the optimal policy characterization. For $0.6 < \alpha \leq \alpha^{\mathrm{C}} = 0.84$, the retailer follows the cutoff partial-disclosure policy. Under this policy, she pools as many low-state predictions as possible with high-state predictions in order to maximize the frequency of high messages, while maintaining just enough credibility for the high message to raise the supplier's posterior to the boundary of the high-production region. Consequently, under the cutoff partial-disclosure policy, π_l is strictly decreasing in α: as α improves, the retailer can send high messages more frequently while still maintaining the minimum credibility needed to induce high-state production.

When α lies between $\alpha^{\mathrm{C}} = 0.84$ and $\alpha^{\mathrm{F}} = 0.91$, the retailer switches to the balanced partial-disclosure policy. In this region, the retailer's prediction has become credible enough to induce reasonably large high-state production, so she balances the credibility of high messages against the frequency of sending such messages. Consequently, π_l becomes strictly increasing in α. Intuitively, as forecast capability improves, the retailer places greater weight on preserving the credibility of high messages and obfuscates her prediction less.

Finally, when $\alpha \geq \alpha^{\mathrm{F}} = 0.91$, maintaining the credibility of high messages dominates, and the retailer returns to full disclosure with $\pi_l^* = 1$. Unlike the full-disclosure region with $\alpha < 0.6$, where the supplier always produces within the low-state demand interval, here a high message induces production within the high-state demand interval.

We now turn to the case of $r = 0.7$ as shown in Figure 3(b), corresponding to Case III of the optimal policy characterization. When $\alpha \leq \alpha^{\mathrm{N}} = 0.94$, the retailer chooses no disclosure, so $\pi_l^* = 0$, because the prior already induces high-state production and noisy forecast

disclosure risks reducing the supplier's posterior. Once α exceeds α^{N}, however, the forecast becomes persuasive enough that revealing high-state predictions creates meaningful benefits to induce greater production quantities. The retailer then switches to balanced partial disclosure and eventually to full disclosure, in which case π_l^* increases with α.

To summarize, when the supplier is reluctant to produce within the high-state demand interval without additional information, due to a low critical ratio, the retailer should fully disclose her forecast when her forecast capability is either sufficiently low or sufficiently high; when forecast capability is moderate, partial disclosure can be optimal. In contrast, when the supplier already has incentives to produce within the high-state demand interval without additional information, no disclosure is preferable unless the retailer's forecast capability is sufficiently high. Once forecast capability becomes sufficiently high, the retailer should disclose more information as forecast capability further improves.

It should be noted that the non-monotonicity of π_l^* only means that the retailer may disclose more or less information about her inherent forecast X, not necessarily about the market state ω. In fact, as forecast capability α improves, even if the retailer obfuscates more of her inherent forecast, she induces a more informative distribution of posterior beliefs about the market state ω, which benefits the supplier.[12]

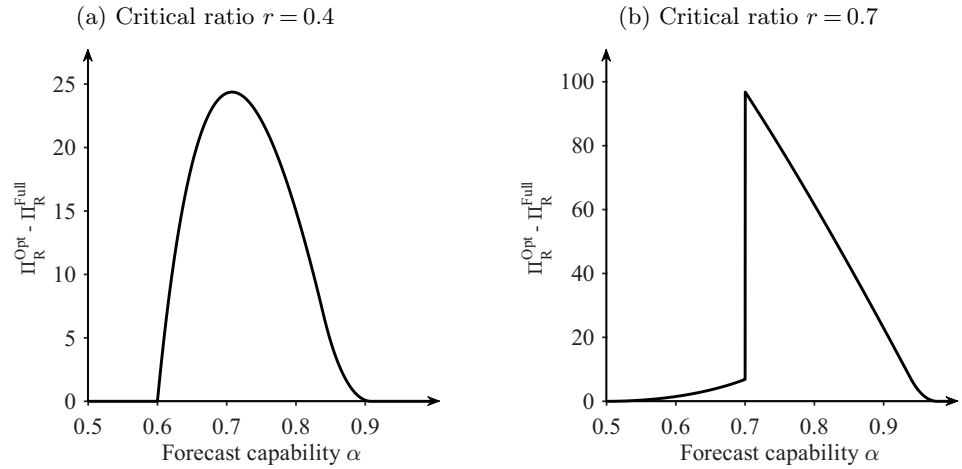
Let Π_R^{Opt} and Π_R^{Full} denote the retailer's expected profits under the optimal disclosure policy and under full disclosure, respectively. Figure 4 further plots the profit improvement relative to full disclosure, $\Pi_R^{\mathrm{Opt}} - \Pi_R^{\mathrm{Full}}$, as forecast capability α varies. For both low and high critical ratios ($r = 0.4$ and $r = 0.7$), the improvement is largest when α is intermediate. For $r = 0.4$, this is immediate, because full disclosure is optimal as α approaches either extreme. For $r = 0.7$, a similar pattern arises because, as α approaches 0.5, the informational value of the retailer's inherent forecast vanishes. Thus, although no disclosure is optimal rather than full disclosure, the resulting profit difference remains small when α is close to 0.5.

4.4. Effect of Demand Range Associated with Each Market State

Under our demand specification, the term $\epsilon \sim U(-\sigma, \sigma)$ captures demand uncertainty conditional on each market state. Holding the retailer's capability to distinguish between the

[12] To see this, whenever a low message is sent under partial or full disclosure, the induced posterior is $\mu(l) = 1 - \alpha$, which decreases with α. The posterior induced by a high message is $\mu(h) = \alpha$ under full disclosure, is fixed at $\mu(h) = 1 - r$ under cutoff partial disclosure, and increases with α under balanced partial disclosure (see Lemma B.1). Hence, as α improves, the induced distribution over μ becomes a mean-preserving spread and provides the supplier with more information about the market state in the sense of Blackwell (1953).

Figure 4 Profit improvement by the optimal disclosure policy relative to full disclosure



Note. $\theta_H = 6$, $\theta_L = 1$, $\sigma = 1$, $p = 1000$, $w = 800$.

high and low market states fixed at α, a larger σ implies a wider demand range, or greater demand uncertainty, within each market state.

One might expect the supplier to always benefit from a smaller σ, since it reduces the inventory mismatch risk conditional on each state. Indeed, it is easy to show that the supplier's profit for any given posterior μ, denoted by $\Pi_S(\mu)$, is decreasing in σ. Thus, an increase in σ has a direct negative effect on the supplier's expected profit. However, a change in σ may also affect how much market-state information the retailer optimally reveals, which in turn influences the distribution of the supplier's posterior beliefs. The following proposition characterizes how σ affects the retailer's optimal disclosure policy and establishes that, because of this indirect disclosure effect, both the supplier's profit and total supply chain profit need not decrease as σ increases.

PROPOSITION 4. *(i) There exist thresholds σ_0, σ_1 and σ_2 such that the following holds: When $1 - \alpha \le r < \frac{1}{2}$ (Case* II*), the optimal policy is full disclosure for $\sigma \ge \sigma_2$, balanced partial disclosure for $\sigma_1 < \sigma < \sigma_2$, and cutoff partial disclosure otherwise; when $r \ge \frac{1}{2}$ (Case* III*), the optimal policy is no disclosure for $\sigma \le \sigma_0$, balanced partial disclosure for $\sigma_0 < \sigma < \sigma_2$, and full disclosure for $\sigma \ge \sigma_2$. (ii) Under balanced partial disclosure, π_l^* is increasing in σ; under all other disclosure policies, π_l^* is constant in σ. (iii) When balanced partial disclosure is optimal for the retailer, an increase in σ raises the supplier's expected profit and lowers the retailer's expected profit, while raising the supply chain's expected profit if and only if p is lower than a threshold $\hat{p}(\sigma)$.*

Generally, as σ increases, the optimal policy shifts from no or partial disclosure toward full disclosure (except in Case I, where full disclosure remains optimal throughout). In particular, under balanced partial disclosure, a larger σ induces the retailer to increase π_l^*, thereby disclosing more information. This suggests that, as the demand range within each state widens, the retailer places greater value on preserving the credibility of the high-demand message.

Part (iii) of Proposition 4 further establishes that, under balanced partial disclosure, the supplier benefits from a larger σ, because the gain from greater information disclosure dominates the direct negative effect of increased within-state uncertainty, whereas the retailer is always worse off as σ increases. Moreover, total supply chain profit may increase with σ when the supplier's profit accounts for a sufficiently large share of total profit, that is, when the retail price p is sufficiently low. Figure 5 presents a numerical illustration when $r = 0.4$: As σ increases, the retailer's optimal policy changes from cutoff partial disclosure, to balanced partial disclosure, and then to full disclosure. The supplier's profit and total supply chain profit both increase with σ within the balanced partial-disclosure region, but decrease once σ moves beyond that region.

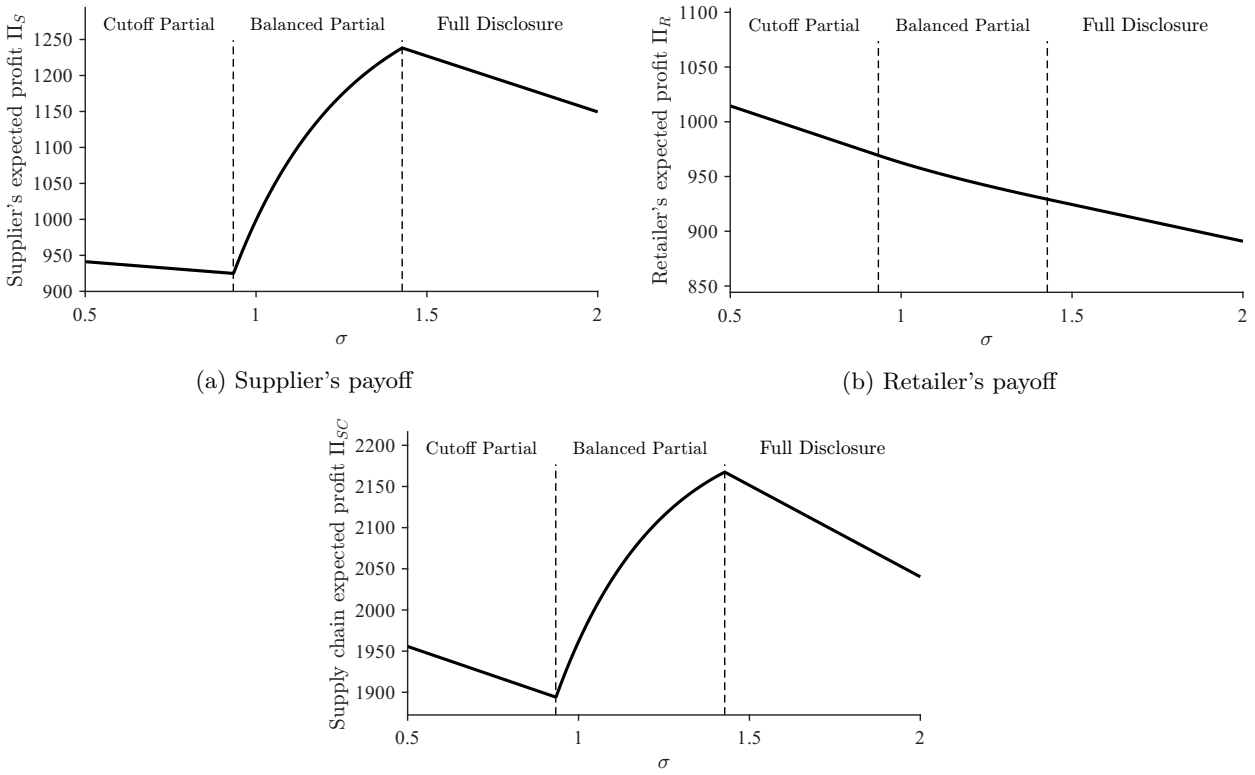
We note that part (iii) of Proposition 4 focuses only on the balanced partial disclosure regime, because the non-monotone effect of σ within that regime arises from the retailer's optimal information disclosure; in other policy regimes, an increase in σ does not influence the retailer's information disclosure and therefore always hurts the supplier. Yet, σ may also have a positive effect on the retailer's profit (and possibly on total supply chain profit) under full disclosure. This occurs because greater demand uncertainty can induce a larger production quantity and higher expected sales when the critical ratio r is sufficiently high in Case I. Since this effect is driven by the supplier's optimal production response rather than the retailer's strategic information disclosure, we present the detailed results in Corollary B.1 in Appendix B.

5. More General Demand Specifications

5.1. Overlapping Uniform Supports

We have thus far assumed that the demand supports under high and low states have nonoverlapping interiors, that is, $\Delta = \theta_H - \theta_L \geq 2\sigma$. We now relax this assumption and consider the case $\Delta < 2\sigma$. As such, the two demand supports, $[\theta_L - \sigma, \theta_L + \sigma]$ and $[\theta_H - \sigma, \theta_H +$

Figure 5 Impact of σ on supply chain profits



(a) Supplier's payoff

(b) Retailer's payoff

(c) Supply chain payoff

Note. $\alpha = 0.85$, $\theta_H = 8$, $\theta_L = 3$, $p = 1000$, $w = 800$, $c = 480$, and hence $r = 0.4$.

$\sigma]$, overlap. Compared with the disjoint-support case, one complication arises because, depending on its posterior belief μ that the market state is high, the supplier's optimal production quantity falls into one of *three* distinct intervals: $[\theta_L - \sigma, \theta_H - \sigma]$, $[\theta_H - \sigma, \theta_L + \sigma]$, or $[\theta_L + \sigma, \theta_H + \sigma]$. We refer to these intervals as the low-, medium-, and high-production intervals, respectively. Let $\mu_1 = 1 - 2r\sigma/\Delta$ and $\mu_2 = 2(1-r)\sigma/\Delta$. It can be shown that the supplier produces within the low-production interval when $\mu \le \mu_1$, within the medium-production interval when $\mu_1 < \mu \le \mu_2$, and within the high-production interval when $\mu > \mu_2$.

Moreover, one can show that the retailer's payoff $\Pi_R(\mu)$ is convex over both $\mu \le \mu_1$ and $\mu_1 < \mu \le \mu_2$, and concave for $\mu > \mu_2$. The slope $\Pi_R'(\mu)$ has a downward kink at $\mu = \mu_1$. This kink occurs between the first two convex pieces, when the posterior belief just enters the medium-production interval and the supplier's production quantity reaches the lower bound of the high-state support, i.e., $Q_\mu^* = \theta_H - \sigma$. As discussed below, the optimal policy may garble the prediction in either direction to keep production at $\theta_H - \sigma$.

We manage to fully characterize the retailer's optimal disclosure policy as presented in Proposition 5.[13] The thresholds used in the proposition, including $r_{\text{ov}}^{\text{C}}(\alpha)$, $r_{\text{ov}}^{\text{DC}}(\alpha)$, $r_{\text{ov}}^{\text{F}}(\alpha)$, and $r_{\text{ov}}^{\text{N}}(\alpha)$, can be derived in closed form and the detailed expressions are provided in the Online Appendices.

PROPOSITION 5. *Suppose that* $0 < \Delta < 2\sigma$. *When* $\alpha = 1$, *full disclosure is optimal. When* $\alpha \in (1/2, 1)$, *the retailer's optimal information policy is characterized as follows.*

- *Case* I*: If* $r < \frac{(1-\alpha)\Delta}{2\sigma}$, *full disclosure is optimal; that is,* $\pi_h^* = \pi_l^* = 1$.
- *Case* II*: If* $\frac{(1-\alpha)\Delta}{2\sigma} \le r < \frac{\Delta}{4\sigma}$, *the optimal policy is characterized as follows.*
  - *(a) If* $r < r_{\text{ov}}^{\text{C}}(\alpha)$, *an upward-cutoff partial disclosure policy is optimal with* $\pi_h^* = 1$ *and* $\pi_l^* = 2 - \frac{1}{\tau_1}$, *where* $\tau_1 = \frac{\alpha - 2r\sigma/\Delta}{2\alpha - 1}$.
  - *(b) If* $r \ge r_{\text{ov}}^{\text{C}}(\alpha)$, *full disclosure is optimal; that is,* $\pi_h^* = \pi_l^* = 1$.
- *Case* III*: If* $r \ge \frac{\Delta}{4\sigma}$, *the optimal policy is characterized as follows.*
  - *(a) If* $r < r_{\text{ov}}^{\text{DC}}(\alpha)$, *a downward-cutoff partial disclosure policy is optimal with* $\pi_h^* = \frac{1-2\tau_1}{1-\tau_1}$ *and* $\pi_l^* = 1$.[14]
  - *(b) If* $r_{\text{ov}}^{\text{DC}}(\alpha) \le r < r_{\text{ov}}^{\text{F}}(\alpha)$, *full disclosure is optimal; that is,* $\pi_h^* = \pi_l^* = 1$.
  - *(c) If* $r_{\text{ov}}^{\text{F}}(\alpha) \le r < r_{\text{ov}}^{\text{N}}(\alpha)$, *a balanced partial disclosure policy is optimal with* $\pi_h^* = 1$ *and* $\pi_l^* = 2 - \frac{1}{\tau_{\text{ov}}^o}$, *where* $\tau_{\text{ov}}^o = \tau^o$ *if* $r \le \frac{\alpha\Delta}{2\sigma}$, *with* τ^o *defined in Equation* (6)*; otherwise,*

$$\tau_{\text{ov}}^o = \left(1 - \frac{\alpha}{1-\alpha} + \frac{2(2\alpha-1)(1-r)\sigma}{(1-\alpha)\sqrt{\alpha\left[4\sigma^2(1-r)^2 - (1-\alpha)^2\Delta^2\right]}}\right)^{-1}. \tag{7}$$

  - *(d) If* $r \ge r_{\text{ov}}^{\text{N}}(\alpha)$, *no disclosure is optimal; that is,* $\pi_h^* = 1$ *and* $\pi_l^* = 0$.

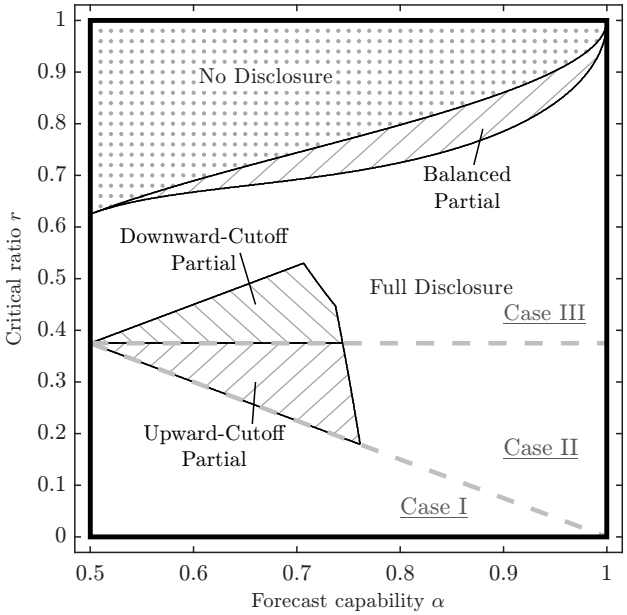
The structure of the optimal policies is illustrated in Figure 6, where the gray dashed lines partition the r–α space into three cases, analogous to the base model. As before, full disclosure is optimal when $\alpha = 1$, but not necessarily when $\alpha < 1$, suggesting that partial disclosure can arise as the optimal policy because of the retailer's imperfect forecast capability.[15] Cases I and II are also akin to the base model. For Case I, where $r < (1-\alpha)\Delta/(2\sigma)$

[13] Because the ordering of the relevant policy boundaries is considerably more involved when expressed in terms of forecast capability α, we characterize how the optimal policy varies with the supplier's critical ratio r.

[14] This policy reduces to no disclosure at $r = \Delta/(4\sigma)$.

[15] The case with perfect forecast capability and overlapping uniform supports is documented in Proposition 5 of the online appendix of Candogan and Gurkan (2025), where full disclosure is shown to be optimal. Here, we show that when forecast capability is imperfect, $\alpha < 1$, more nuanced partial- or no-disclosure policies can be optimal.

Figure 6 **Structure of the optimal policy with overlapping supports $\Delta < 2\sigma$**



Note. $\Delta = 1.5$, $\sigma = 1$. Positive-slope hatching represents upward-cutoff or balanced partial disclosure, while negative-slope hatching represents downward-cutoff partial disclosure. The gray dashed boundaries separate Cases I–III in Proposition 5.

or equivalently $\alpha < \mu_1$, both the supplier's critical ratio r and the retailer's forecast capability are low such that all the inducible posteriors lie below μ_1, so full disclosure is optimal as explained in § 4.2.2. For Case II, the condition $(1-\alpha)\Delta/(2\sigma) \le r < \Delta/(4\sigma)$ is equivalent to $1/2 < \mu_1 \le \alpha$. Thus, the supplier can be induced to produce beyond the low-production interval under sufficiently favorable posterior beliefs. However, without additional information, the supplier produces within the low-production interval because the prior belief lies below μ_1. In this case, the optimal policy follows the same pattern as in the base model. As shown in Figure 6, when the critical ratio r (or the forecast capability α) is low, the optimal policy exaggerates low predictions as high with probability $1-\pi_l^*$, which induces exactly a posterior $\mu(h) = \mu_1$ under a high message. This preserves the supplier's production response at exactly $Q^*_{\mu(h)} = \theta_H - \sigma$ while maximizing the high-message frequency. We refer to this policy as the *upward-cutoff partial disclosure* policy. By contrast, when r (or α) is high, full disclosure becomes optimal.

Case III corresponds to the setting in which the prior belief already induces a production quantity beyond the low-production interval, i.e., $Q^*_{1/2} > \theta_H - \sigma$. This case becomes more involved than that in the base model. Interestingly, when $r < r_{\text{ov}}^{\text{DC}}(\alpha)$, the optimal policy

may downplay a high-state prediction by sending a low message while always sending a low message following a low-state prediction, i.e., $\pi_h^* < 1$ and $\pi_l^* = 1$, as characterized in subcase (a) of Case III. This stands in sharp contrast to the base model, in which the retailer never downplays high-state predictions. Under such a policy, which we refer to as *downward-cutoff partial disclosure*, the retailer pools some high-state predictions into the low message so that a low message induces a posterior exactly at $\mu(l) = \mu_1$. Doing so preserves the supplier's production quantity at $Q^*_{\mu(l)} = \theta_H - \sigma$, while reducing the frequency of the high message. This policy is optimal when r is relatively low, because the supplier's incentive to produce a large quantity following a high message is weak, making the loss from reducing the high-message frequency relatively small.

As r increases beyond $r_{\text{ov}}^{\text{DC}}(\alpha)$, preserving the quantity $Q^*_{\mu(l)} = \theta_H - \sigma$ while sacrificing high-message frequency becomes less valuable since the higher r makes the production response induced by high messages sufficiently valuable. Thus, when $r_{\text{ov}}^{\text{DC}}(\alpha) \le r < r_{\text{ov}}^{\text{F}}(\alpha)$, full disclosure becomes optimal. When $r_{\text{ov}}^{\text{F}}(\alpha) \le r < r_{\text{ov}}^{\text{N}}(\alpha)$, increasing the high-message frequency through exaggerating some low predictions becomes more valuable, so the optimal policy becomes balanced partial. As discussed in the base model, this policy induces a posterior $\mu(h)$ lying within the interior of the high-production interval $[\mu_2, 1]$, balancing the credibility-frequency trade-off. When $r \ge r_{\text{ov}}^{\text{N}}(\alpha)$, no disclosure becomes optimal, as the prior already induces high production, and the limited benefit of a favorable message no longer offsets the production reduction caused by an unfavorable message.

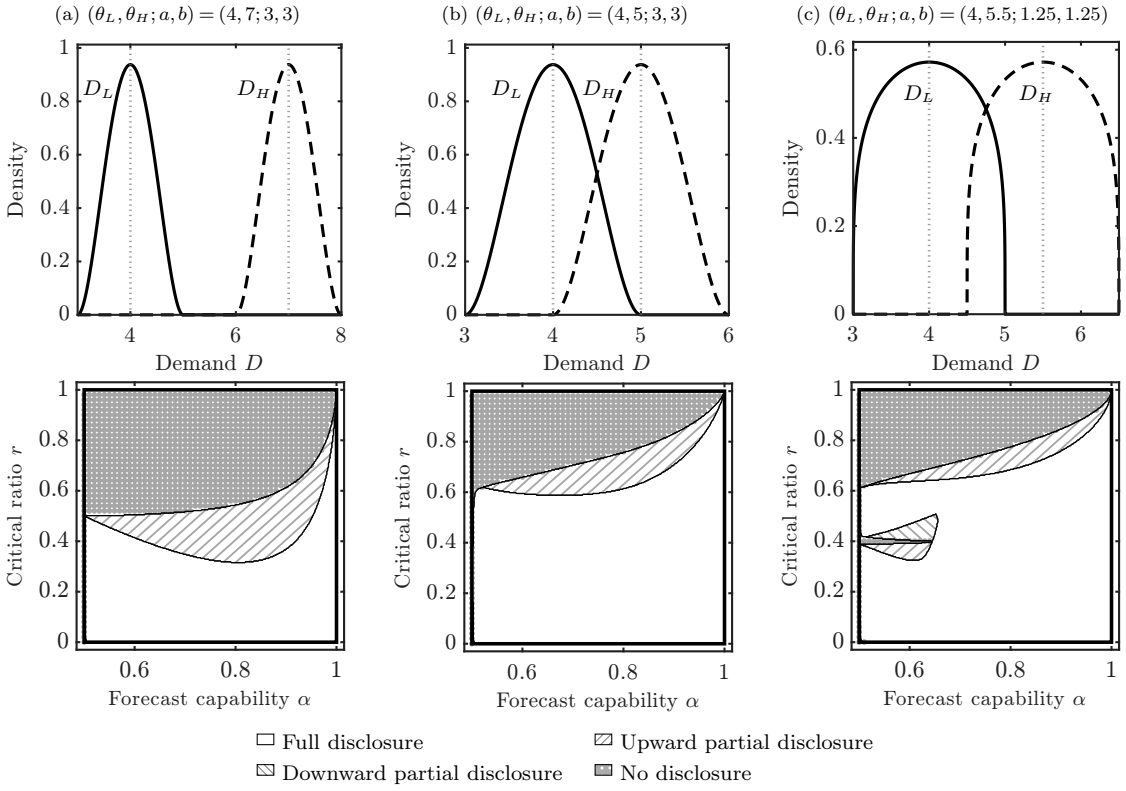
In summary, with overlapping supports, imperfect forecast capability still gives rise to a nontrivial credibility–frequency tradeoff, such that partial or no disclosure can be optimal. Moreover, overlap can make either upward- or downward-cutoff partial disclosure optimal, while both policies serve to keep production within the high-state support.

5.2. Non-uniform Demand

We numerically examine more general demand distributions as a robustness check for our analytical results derived under uniform demand. Specifically, we conduct two sets of numerical studies in which demand follows state-dependent location-scale Beta and Gamma distributions, respectively.

5.2.1. Location-Scale Beta Distribution We assume that conditional on the market state $\omega \in \{L, H\}$, demand follows a location-scale Beta distribution: $D_\omega = \theta_\omega - \frac{2\sigma a}{a+b} + 2\sigma Y$, and $Y \sim \text{Beta}(a, b)$, where $a, b > 0$ are the shape parameters. We choose this specification because (1) $\mathbb{E}[D_\omega] = \theta_\omega$, (2) the two demand distributions have the same shape and dispersion, while differing only by the location shift $\Delta = \theta_H - \theta_L$, and (3) their supports are nonoverlapping when $\Delta \geq 2\sigma$ and overlapping when $\Delta < 2\sigma$. Furthermore, the baseline uniform distribution is nested as the special case $a = b = 1$.

Figure 7 Optimal policy under Beta demand $D_\omega = \theta_\omega - \frac{2a}{a+b} + 2Y$ where $Y \sim \text{Beta}(a, b)$ and $\omega \in \{L, H\}$



Note. For all distributions, we set $\sigma = 1$.

We consider three different distributional configurations with $\sigma = 1$. The remaining parameter values and the corresponding density functions are shown in Cases (a)–(c) of Figure 7. For each configuration, we numerically solve for the optimal disclosure policy across different combinations of r and α and plot the resulting optimal-policy regions in the bottom panels of Figure 7. The numerical results confirm that our analytical findings are robust. In Cases (a) and (b), the optimal policies retain a similar structure as in the base

model. Across the three configurations examined, upward partial disclosure predominantly takes the balanced form. In Case (c), the overlapping region carries relatively high probability mass from both demand states, as the high-state density rises sharply upon entering the overlap while the low-state density remains substantial. Consequently, the optimal policy exhibits a structure similar to that in the overlapping-support setting analyzed in §5.1. In particular, when r is around 0.4 and α is relatively small, a disconnected region in which partial or no disclosure is optimal emerges. Within this region, downward partial disclosure can be optimal. These observations suggest that downward partial disclosure arises when probability mass is highly concentrated in the overlapping region.

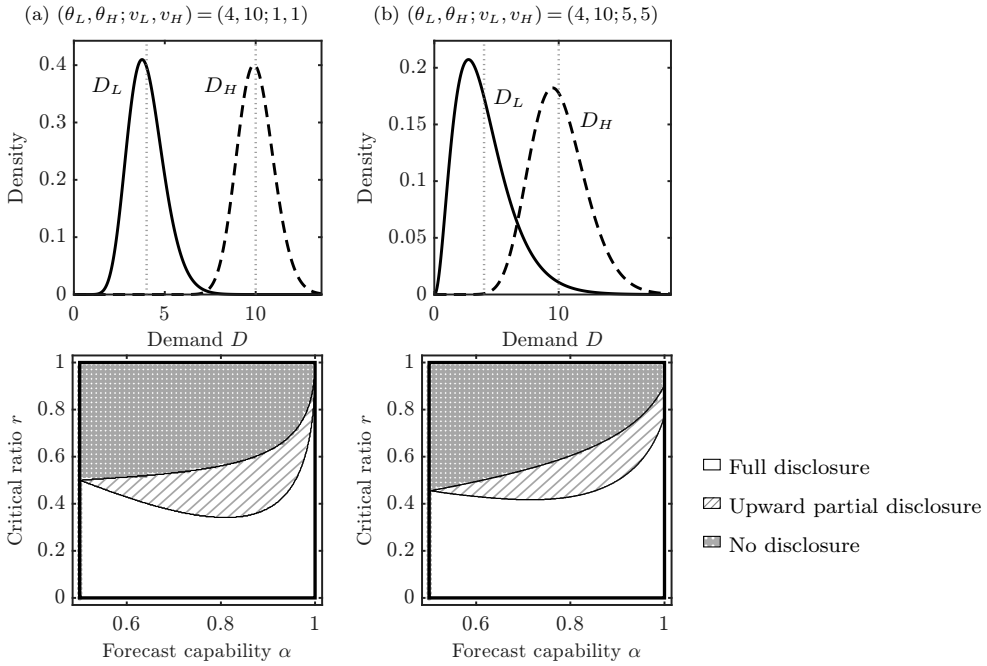
5.2.2. Gamma Distribution Demand Next, we consider Gamma-distributed demand, for which demand supports become unbounded. Specifically, conditional on the market state $\omega \in \{L, H\}$, demand $D_\omega \sim \text{Gamma}\,(\beta_\omega, \gamma_\omega)$, where β_ω and γ_ω denote the shape and scale parameters, respectively. Accordingly, D_ω has mean $\mathbb{E}[D_\omega] = \beta_\omega \gamma_\omega$ and variance $\text{Var}(D_\omega) = \beta_\omega \gamma_\omega^2$. As reported in Figure 8, we examine two different cases and choose the shape and scale parameters so that D_L and D_H have different means and the same variance in each case: $v_L = v_H = 1$ in Case (a) and $v_L = v_H = 5$ in Case (b). The numerical results again confirm that the optimal policy exhibits a structure similar to that in our base model.[16] Unlike bounded-support distributions, Gamma demand distributions share the same unbounded support over the nonnegative real line. In the two configurations examined, whenever partial disclosure is optimal, it takes the balanced form.

6. Conclusion

Building on recent literature on information design in supply chains, we study how a retailer should share its market-state forecast with an upstream supplier to persuade the supplier to choose higher production levels. We depart from this literature by allowing the retailer's forecast capability to be imperfect. Our findings advance the current understanding of information-sharing incentives in supply chains. In particular, our results imply that partial or no disclosure may result from a retailer's imperfect forecast capability. Under state-dependent uniform demand distributions that differ only by a location shift, we explicitly

[16] Under state-dependent Gamma distributions, the demand density may take different shapes across market states. Thus, even when $\alpha = 1$, full disclosure need not be optimal if the density shapes differ sufficiently across states; see, for example, Case (b). Nevertheless, the role of forecast capability α in shaping the optimal disclosure policy remains similar to that identified in our earlier analysis.

Figure 8 **Optimal policy under Gamma demand with mean θ_ω and variance v_ω where $\omega \in \{H, L\}$**



Note. Cases (a) and (b) use $(\beta_L, \gamma_L; \beta_H, \gamma_H) = (16, 0.25; 100, 0.1)$ and $(3.2, 1.25; 20, 0.5)$, respectively. These parameterizations correspond to $(\theta_L, \theta_H; v_L, v_H) = (4, 10; 1, 1)$ and $(4, 10; 5, 5)$, respectively, where $\theta_\omega = \mathbb{E}[D_\omega]$ and $v_\omega = \mathrm{Var}(D_\omega)$.

characterize the retailer's optimal information policy. Our characterization prescribes that retailers tailor their information-sharing policies to both their forecast capability and their suppliers' production incentives.

Our analysis suggests several directions for future research. First, our model assumes that the retailer's forecast accuracy is exogenous. Allowing the retailer to invest in forecast capability before choosing a disclosure policy would enable the joint study of forecasting investment and information disclosure. Second, we focus on a single retailer and a single supplier. Future research could consider multiple competing retailers or a retailer sourcing from multiple suppliers. Finally, we study information design under a fixed wholesale-price contract. Exploring richer contractual arrangements or jointly designing information disclosure and contractual terms would also be valuable.

References

Alizamir S, de Véricourt F, Wang S (2020) Warning against recurring risks: An information design approach. *Management Science* 66(10):4612–4629.

Best J, Quigley D (2024) Persuasion for the long run. *Journal of Political Economy* 132(5):1740–1791.

Bimpikis K, Papanastasiou Y, Zhang W (2024) Information provision in two-sided platforms: Optimizing for supply. *Management Science* 70(7):4533–4547.

Blackwell D (1953) Equivalent comparisons of experiments. *The Annals of Mathematical Statistics* 24(2):265–272.

Boleslavsky R, Cotton CS, Gurnani H (2017) Demonstrations and price competition in new product release. *Management Science* 63(6):2016–2026.

Boyaci T, Chakraborty S, Gurkan H (2024) Persuading skeptics and fans in the presence of additional information. *Production and Operations Management* 33(5):1142–1154.

Cachon GP, Lariviere MA (2001) Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science* 47(5):629–646.

Candogan O (2020) Information design in operations. *Pushing the Boundaries: Frontiers in Impactful OR/OM Research*, 176–201 (INFORMS).

Candogan O, Gurkan H (2025) The value of information design in supply chain management. *Management Science* 71(8):6545–6558.

Candogan O, Gurkan H, Ismayilov R (2026) Information design in supply chains with priority suppliers. *Available at SSRN 4890943* .

Chen Y, Ma G, He QC, Chen YJ, Shen ZJM (2026) Strategic data analytics under supplier encroachment. *Manufacturing & Service Operations Management* 28(4):1324–1338.

Chen YJ, Xiao W (2012) Impact of reseller's forecasting accuracy on channel member performance. *Production and Operations Management* 21(6):1075–1089.

Chu LY, Shamir N, Shin H (2017) Strategic communication for capacity alignment with pricing in a supply chain. *Management Science* 63(12):4366–4388.

De Véricourt F, Gurkan H, Wang S (2021) Informing the public about a pandemic. *Management Science* 67(10):6350–6357.

Drakopoulos K, Jain S, Randhawa R (2021) Persuading customers to buy early: The value of personalized information provisioning. *Management Science* 67(2):828–853.

Farahani MH, Dawande M, Janakiraman G, Wang S (2025) Avoiding fields on fire: Information dissemination policies for environmentally safe crop-residue management. *Management Science* 71(8):6683–6706.

Gümüş M (2014) With or without forecast sharing: Competition and credibility under information asymmetry. *Production and Operations Management* 23(10):1732–1747.

Ha AY, Tong S (2008) Contracting and information sharing under supply chain competition. *Management Science* 54(4):701–715.

Hamada A (2025) Amazon demand forecasting for vendors. Reason Automation, Guides & How-to, URL `https://www.reasonautomation.com/resource-library/amazon-demand-forecasting-for-vendors`, accessed: June 18, 2026.

Huang S, Guan X, Chen YJ (2018) Retailer information sharing with supplier encroachment. *Production and Operations Management* 27(6):1133–1147.

Iyer G, Narasimhan C, Niraj R (2007) Information and inventory in distribution channels. *Management Science* 53(10):1551–1561.

Jain A (2022) Sharing demand information with retailer under upstream competition. *Management Science* 68(7):4983–5001.

Jain A, Seshadri S, Sohoni M (2011) Differential pricing for information sharing under competition. *Production and Operations Management* 20(2):235–252.

Jiang B, Tian L, Xu Y, Zhang F (2016) To share or not to share: Demand forecast sharing in a distribution channel. *Marketing Science* 35(5):800–809.

Kamenica E, Gentzkow M (2011) Bayesian persuasion. *American Economic Review* 101(6):2590–2615.

Küçükgül C, Özer Ö, Wang S (2022) Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Management Science* 68(7):4899–4918.

Liu Z, Zhang DJ, Zhang F (2021) Information sharing on retail platforms. *Manufacturing & Service Operations Management* 23(3):606–619.

Lu T, Tomlin B (2025) Augmenting the operations manager with a prediction machine. *Tuck School of Business Working Paper* .

Lu T, Tomlin B (2026) Credible cheap-talk communication of private demand information on both the forecast average and accuracy. *Management Science* 72(2):989–1006.

Özer Ö, Wei W (2006) Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Science* 52(8):1238–1257.

Özer Ö, Zheng Y, Chen KY (2011) Trust in forecast information sharing. *Management Science* 57(6):1111–1137.

Qi Y, Beamish PW, Li X, Gong K (2021) Supply chain collaboration at JD.com. Ivey Publishing Case Study, product No. W21286; Ivey case no. 9B21D009.

Shamir N, Shin H (2016) Public forecast information sharing in a market with competing supply chains. *Management Science* 62(10):2994–3022.

Shang W, Ha AY, Tong S (2016) Information sharing in a supply chain with a common retailer. *Management Science* 62(1):245–263.

Taylor TA, Xiao W (2010) Does a manufacturer benefit from selling to a better-forecasting retailer? *Management Science* 56(9):1584–1598.

Wang Y, Ha AY, Tong S (2022) Sharing manufacturer's demand information in a supply chain with price and service effort competition. *Manufacturing & Service Operations Management* 24(3):1698–1713.

Xu A, Chen Z, He QC, Chen YJ (2026) Strategic communications with socializing agents under unknown public health threats. *Manufacturing & Service Operations Management* 28(1):42–56.

Yang W, Yao Y, Zhou P (2026) Strategic disinformation generation and detection. *Management Science* .

Online Appendices for "Persuading the Supplier with Imperfect Forecast Capability"

For notational convenience, throughout these appendices we write $\mu_s = \mu(s)$ and $\tau_s = \Pr(X = 1 \mid s)$ for $s \in \{l, h\}$.

A Proofs of Results in the Main Text

Proof of Lemma 1. Fix a posterior belief $\mu \in [0, 1]$. The c.d.f. of demand conditional on μ is

$$F_D(x;\mu) = \begin{cases} 0, & x < \theta_L - \sigma, \\ (1-\mu)\dfrac{x - (\theta_L - \sigma)}{2\sigma}, & \theta_L - \sigma \le x < \theta_L + \sigma, \\ 1 - \mu, & \theta_L + \sigma \le x < \theta_H - \sigma, \\ 1 - \mu + \mu\dfrac{x - (\theta_H - \sigma)}{2\sigma}, & \theta_H - \sigma \le x \le \theta_H + \sigma, \\ 1, & x > \theta_H + \sigma. \end{cases}$$

By the standard newsvendor critical-fractile condition, an optimal production quantity satisfies $F_D(Q;\mu) = r$. At $\mu = 1 - r$, every quantity in $[\theta_L + \sigma, \theta_H - \sigma]$ is optimal, and the stated tie-breaking rule selects $Q^*_{1-r} = \theta_H - \sigma$. Solving on the two branches yields (5).

For $\mu \neq 1 - r$, differentiating (5) gives

$$\frac{dQ^*_\mu}{d\mu} = \begin{cases} \dfrac{2\sigma r}{(1-\mu)^2}, & \mu < 1 - r, \\ \dfrac{2\sigma(1-r)}{\mu^2}, & \mu > 1 - r, \end{cases} \qquad \frac{d^2Q^*_\mu}{d\mu^2} = \begin{cases} \dfrac{4\sigma r}{(1-\mu)^3}, & \mu < 1 - r, \\ -\dfrac{4\sigma(1-r)}{\mu^3}, & \mu > 1 - r. \end{cases}$$

At the cutoff, $Q^*_{1-r} - \lim_{\mu \to (1-r)^-} Q^*_\mu = \Delta - 2\sigma \ge 0$. Thus, Q^*_μ jumps upward when $\Delta > 2\sigma$ and is continuous when $\Delta = 2\sigma$. Together with the derivative signs and continuity of the upper branch at its left endpoint, this establishes the stated monotonicity and curvature properties. Q.E.D.

Proof of Lemma 2. By Lemma 1, when $\mu < 1 - r$, Q^*_μ lies in the low-demand support. Hence, $\mathbb{E}[\min(\theta_H + \epsilon, Q^*_\mu)] = Q^*_\mu$ and $\mathbb{E}[\min(\theta_L + \epsilon, Q^*_\mu)] = \theta_L - (\theta_L + \sigma - Q^*_\mu)^2/(4\sigma)$. Substituting Q^*_μ from Lemma 1 into the retailer's profit gives the first branch in Lemma 2.

When $\mu \geq 1-r$, Q_μ^* lies in the high-demand support. Thus, $\mathbb{E}[\min(\theta_L+\epsilon, Q_\mu^*)] = \theta_L$ and $\mathbb{E}[\min(\theta_H+\epsilon, Q_\mu^*)] = \theta_H - (\theta_H+\sigma-Q_\mu^*)^2/(4\sigma)$, which similarly gives the second branch in Lemma 2.

For $\mu \neq 1-r$, differentiating the two branches gives

$$\Pi_R'(\mu) = \begin{cases} (p-w)\sigma\dfrac{r(2-r)}{(1-\mu)^2}, & \mu < 1-r, \\ (p-w)\left[\Delta + \dfrac{\sigma(1-r)^2}{\mu^2}\right], & \mu > 1-r, \end{cases}$$

which is strictly positive. Similarly,

$$\Pi_R''(\mu) = \begin{cases} 2(p-w)\sigma\dfrac{r(2-r)}{(1-\mu)^3}, & \mu < 1-r, \\ -2(p-w)\sigma\dfrac{(1-r)^2}{\mu^3}, & \mu > 1-r. \end{cases}$$

Hence, $\Pi_R(\mu)$ is strictly convex for $\mu < 1-r$ and strictly concave for $\mu > 1-r$.

Finally, at $\mu = 1-r$,

$$\Pi_R(1-r) - \lim_{\mu \to (1-r)^-} \Pi_R(\mu) = (p-w)(1-r)(\Delta - 2\sigma) \geq 0.$$

Thus, $\Pi_R(\mu)$ has an upward jump at $\mu = 1-r$ when $\Delta > 2\sigma$, and is continuous there when $\Delta = 2\sigma$. The derivative signs, this nonnegative jump, and continuity of the upper branch at its left endpoint establish the stated monotonicity and curvature properties. Q.E.D.

The following equivalent formulation in Corollary A.1 will be useful for the subsequent analysis.

COROLLARY A.1. *The retailer's information-design problem of maximizing expected profit is equivalent to minimizing expected shortage cost.*

Proof of Corollary A.1. Under policy π, using $\min(D,Q) = D - (D-Q)^+$, the retailer's expected profit can be written as $(p-w)\mathbb{E}[D] - (p-w)\mathbb{E}\left[(D-Q_\mu^*)^+\right]$. Since $\mathbb{E}[D]$ is independent of π, the result follows. Q.E.D.

Proof of Proposition 2. In this proof, we first characterize the shape of the retailer's expected shortage cost in the posterior index τ, which allows us to reduce the information-design problem to choosing a split between $\tau_l = 0$ and a high posterior index τ_h. We then minimize the expected cost over τ_h, use the resulting boundary conditions to identify the thresholds α^{C}, α^{N}, and α^{F}, and characterize the optimal disclosure policy in each case.

We now provide the details. By Corollary A.1, we work with the retailer's expected shortage cost. Recall that $\mu(\tau) = 1-\alpha+(2\alpha-1)\tau$ for $\tau \in [0,1]$. Thus, $\tau = 0$, $1/2$, and 1 correspond to posterior beliefs $1-\alpha$, $1/2$, and α, respectively. In the τ-representation, the expected shortage cost is

$$V(\tau) = \begin{cases} \Delta\mu(\tau) + \dfrac{\sigma\big((1-r)^2 - \mu(\tau)\big)}{1-\mu(\tau)}, & \tau < \tau^\dagger, \\ \dfrac{\sigma(1-r)^2}{\mu(\tau)}, & \tau \geq \tau^\dagger, \end{cases}$$

where $\tau^\dagger = (\alpha - r)/(2\alpha - 1)$. Let $V_1(\tau)$ and $V_2(\tau)$ denote this cost on the branches $\tau < \tau^\dagger$ and $\tau \geq \tau^\dagger$, respectively.

Direct differentiation gives $V_1''(\tau) = 2\sigma(2\alpha-1)^2[(1-r)^2-1]/[1-\mu(\tau)]^3 < 0$ and $V_2''(\tau) = 2\sigma(1-r)^2(2\alpha-1)^2/\mu(\tau)^3 > 0$. Hence V_1 is strictly concave and V_2 is strictly convex. At the cutoff, $V_1((\tau^\dagger)^-) - V_2(\tau^\dagger) = (1-r)(\Delta - 2\sigma) \geq 0$.

Suppose first that $0 < \tau^\dagger \leq 1$. For any $\tau \in (0, \tau^\dagger)$, the strict concavity of V_1 and the nonnegative gap at $\tau^\dagger$ imply

$$V_1(\tau) > \left(1 - \frac{\tau}{\tau^\dagger}\right) V_1(0) + \frac{\tau}{\tau^\dagger} V_1((\tau^\dagger)^-) \geq \left(1 - \frac{\tau}{\tau^\dagger}\right) V_1(0) + \frac{\tau}{\tau^\dagger} V_2(\tau^\dagger).$$

The mixture on the right-hand side has the same mean posterior index τ but a strictly lower expected cost. Hence no interior point of the V_1 branch can belong to an optimal posterior support. Moreover, because V_2 is convex, any distribution supported on multiple points of V_2 can be pooled at its probability-weighted mean without increasing expected cost. Therefore, whenever $0 < \tau^\dagger \leq 1$, any informative optimal policy can be represented by a posterior distribution supported on $\tau_l = 0$ and a single $\tau_h \in [\max\{\tau^\dagger, 1/2\}, 1]$. The configurations $\tau^\dagger > 1$ and $\tau^\dagger \leq 0$, in which the feasible interval lies on a single branch, are handled directly below.

Let $P := V_1(0) = \Delta(1-\alpha) + \sigma[(1-r)^2 - (1-\alpha)]/\alpha$. For a split supported on 0 and $\tau_h \in [1/2, 1]$, Bayes plausibility implies that the weights on τ_h and 0 are $1/(2\tau_h)$ and $1 - 1/(2\tau_h)$, respectively. This split is implemented by $\pi_h = 1$ and $\pi_l = 2 - 1/\tau_h$. Its expected shortage cost and derivative can be written as

$$\begin{aligned} &L(\tau_h) = \left(1 - \frac{1}{2\tau_h}\right) P + \frac{V_2(\tau_h)}{2\tau_h}, \quad L'(\tau_h) = \frac{\Psi(\tau_h)}{2\tau_h^2}, \\ &\Psi(\tau_h) = \tau_h V_2'(\tau_h) - V_2(\tau_h) + P, \quad \Psi'(\tau_h) = \tau_h V_2''(\tau_h) > 0. \end{aligned} \tag{A.1}$$

Thus, Ψ is strictly increasing. If its signs at the endpoints of the relevant interval differ, L has a unique interior minimizer characterized by $\Psi(\tau_h) = 0$. Solving this equation gives the positive root τ^o in (6).

For compactness, define G_{F}, G_{C}, and G_{N} as follows. Direct substitution into (A.1) yields

$$\begin{aligned} G_{\mathrm{F}}(\alpha) &= \Delta\alpha^3 - (\Delta+\sigma)\alpha^2 + \sigma[1+2(1-r)^2]\alpha - \sigma(1-r)^2, \\ G_{\mathrm{C}}(\alpha) &= (\Delta+\sigma)\alpha^2 - (\Delta+2r\sigma)\alpha + \sigma r(2-r), \\ G_{\mathrm{N}}(\alpha) &= [\Delta+4\sigma(1-r)^2]\alpha^2 - (\Delta+\sigma)\alpha + \sigma r(2-r), \\ \Psi(1) &= -\frac{G_{\mathrm{F}}(\alpha)}{\alpha^2}, \qquad \Psi(\tau^\dagger) = -\frac{G_{\mathrm{C}}(\alpha)}{\alpha}, \qquad \Psi\left(\frac{1}{2}\right) = -\frac{G_{\mathrm{N}}(\alpha)}{\alpha}. \end{aligned} \tag{A.2}$$

These polynomials also establish the direction and uniqueness of the threshold comparisons. Suppose first that $\Delta > 2\sigma$. Then $G_{\mathrm{F}}(1/2) = (2\sigma-\Delta)/8 < 0$ and $G_{\mathrm{F}}(1) = \sigma(1-r)^2 > 0$. Setting $x = 2\alpha - 1$, the coefficients of $8G_{\mathrm{F}}((1+x)/2)$, in ascending powers of x, are $2\sigma-\Delta$, $-\Delta+8\sigma(1-r)^2$, $\Delta-2\sigma$, and Δ. They exhibit exactly one sign change regardless of the sign of the second coefficient. Descartes' rule of signs therefore implies that G_{F} has exactly one root α^{F} in $(1/2,1)$, with $G_{\mathrm{F}}(\alpha) < 0$ below it and $G_{\mathrm{F}}(\alpha) > 0$ above it. Similarly, $G_{\mathrm{C}}(1/2) = [-\Delta+\sigma+4r(1-r)\sigma]/4 < 0$, $G_{\mathrm{C}}(1) = \sigma(1-r^2) > 0$, $G_{\mathrm{N}}(1/2) = (2\sigma-\Delta)/4 < 0$, and $G_{\mathrm{N}}(1) = 3\sigma(1-r)^2 > 0$. Since G_{C} and G_{N} are upward-opening quadratics, each has a unique root in $(1/2,1)$, namely α^{C} and α^{N}, respectively.

When $\Delta = 2\sigma$, direct factorization gives $G_{\mathrm{F}}(\alpha) = \sigma(2\alpha-1)[\alpha^2-\alpha+(1-r)^2]$, $G_{\mathrm{C}}(\alpha) = \sigma(\alpha-r)(3\alpha+r-2)$, and $G_{\mathrm{N}}(\alpha) = \sigma(2\alpha-1)\{[1+2(1-r)^2]\alpha - [1-(1-r)^2]\}$. In this boundary case, the threshold formulas stated above simplify to $\alpha^{\mathrm{C}} = (2-r)/3$ in Case II, where $r < 1/2$, and $\alpha^{\mathrm{N}} = [1-(1-r)^2]/[1+2(1-r)^2]$ in Case III, where $r \geq 1/2$. Each expression selects the larger root of the corresponding quadratic. Moreover, $\alpha^{\mathrm{F}} = 1/2$ for $r \leq 1/2$, whereas for $r > 1/2$, it is the unique root of G_{F} in $(1/2,1)$. Together with the preceding analysis for $\Delta > 2\sigma$, these factorizations establish the signs needed below: within $(1/2,1)$, each relevant polynomial is negative below its corresponding threshold and positive above it. When the threshold equals $1/2$, the polynomial is positive throughout $(1/2,1)$.

We now prove the result case by case.

Case I: $r < 1-\alpha$. Here $\tau^\dagger > 1$, so $V = V_1$ throughout $[0,1]$. Since V_1 is concave, the cost-minimizing Bayes-plausible posterior distribution places equal mass on the two endpoints 0 and 1. This is full disclosure.

Case II: $1-\alpha \le r < 1/2$. In this case, $\tau^\dagger \in (1/2, 1]$, so the relevant choice in (A.1) is $\tau_h \in [\tau^\dagger, 1]$. Suppose first that $\alpha^{\mathrm{C}} < 1-r$. Then $G_{\mathrm{C}}(1-r) > 0$. Since $G_{\mathrm{F}}(1-r) = (1-r)G_{\mathrm{C}}(1-r) > 0$, the sign characterization above implies $\alpha^{\mathrm{F}} < 1-r$. Every α in Case II therefore satisfies $\alpha \ge 1-r > \alpha^{\mathrm{F}}$, so $\Psi(1) < 0$ and full disclosure is optimal.

It remains to consider $\alpha^{\mathrm{C}} \ge 1-r$, so that the threshold point lies in Case II. At $\alpha = \alpha^{\mathrm{C}}$, we have $\Psi(\tau^\dagger) = 0$ and $\tau^\dagger \le 1$. Since Ψ is strictly increasing, $\Psi(1) \ge 0$. By (A.2), this implies $G_{\mathrm{F}}(\alpha^{\mathrm{C}}) \le 0$, and hence $\alpha^{\mathrm{C}} \le \alpha^{\mathrm{F}}$. The inequality is strict when $\alpha^{\mathrm{C}} > 1-r$; when $\alpha^{\mathrm{C}} = 1-r$, $\tau^\dagger = 1$ and $\alpha^{\mathrm{C}} = \alpha^{\mathrm{F}} = 1-r$.

(a) If $\alpha \ge \alpha^{\mathrm{F}}$, then $\Psi(1) \le 0$ by (A.2). Since Ψ is increasing, $\Psi(\tau_h) \le 0$ throughout $[\tau^\dagger, 1]$, so L is decreasing on this interval. Hence $\tau_h = 1$, which implements full disclosure.

(b) If $\alpha^{\mathrm{C}} < \alpha < \alpha^{\mathrm{F}}$, then $\Psi(\tau^\dagger) < 0 < \Psi(1)$. Since Ψ is strictly increasing, L is minimized at the unique solution $\tau_h = \tau^o \in (\tau^\dagger, 1)$ of $\Psi(\tau_h) = 0$, which implements $\pi_h = 1$ and $\pi_l = 2 - 1/\tau^o$.

(c) Otherwise, $\alpha < \alpha^{\mathrm{F}}$ and $\alpha \le \alpha^{\mathrm{C}}$, so $\Psi(\tau^\dagger) \ge 0$. Since Ψ is increasing, $\Psi(\tau_h) \ge 0$ throughout $[\tau^\dagger, 1]$, so L is increasing on this interval. Hence its minimizer is $\tau_h = \tau^\dagger$, which implements $\pi_h = 1$ and $\pi_l = 2 - 1/\tau^\dagger$. This completes Case II.

Case III: $r \ge 1/2$. If $\Delta = 2\sigma$ and $r = 1/2$, then $\tau^\dagger = 1/2$, and the factorization above gives $G_{\mathrm{F}}(\alpha) > 0$ for every $\alpha \in (1/2, 1)$. Hence $\Psi(1) < 0$, so the strict increase of Ψ implies that L is decreasing on $[1/2, 1]$, and full disclosure is optimal. For all remaining parameter configurations in Case III, $\alpha^{\mathrm{N}} > 1/2$. We first record the ordering of the two relevant thresholds. Direct substitution gives $G_{\mathrm{N}}(r) = r(r-1)[\Delta - \sigma - 4r(1-r)\sigma] \le 0$, where the inequality follows from $r \in (0, 1)$ and $\Delta \ge 2\sigma$. Hence $r \le \alpha^{\mathrm{N}}$. Moreover, at $\alpha = \alpha^{\mathrm{N}}$, $\Psi(1/2) = 0$. Since Ψ is strictly increasing, $\Psi(1) > 0$, so $G_{\mathrm{F}}(\alpha^{\mathrm{N}}) < 0$, which implies $\alpha^{\mathrm{N}} < \alpha^{\mathrm{F}}$.

If $\alpha \le r$, then $\tau^\dagger \le 0$, so $V = V_2$ throughout $[0, 1]$. Since V_2 is convex, Jensen's inequality implies that no disclosure, which keeps the posterior at $\tau = 1/2$, minimizes expected cost.

It remains to consider $\alpha > r$. Then $\tau^\dagger \in (0, 1/2]$, so the relevant choice in (A.1) is $\tau_h \in [1/2, 1]$.

(a) If $\alpha \ge \alpha^{\mathrm{F}}$, then $\Psi(1) \le 0$. Since Ψ is increasing, L is decreasing on $[1/2, 1]$. Hence $\tau_h = 1$, and full disclosure is optimal.

(b) If $\alpha^{\mathrm{N}} < \alpha < \alpha^{\mathrm{F}}$, then $\Psi(1/2) < 0 < \Psi(1)$. Hence L has a unique minimizer $\tau_h = \tau^o \in (1/2, 1)$, which implements the balanced partial-disclosure policy stated in the proposition.

(c) If $r < \alpha \leq \alpha^{\mathrm{N}}$, then $\Psi(1/2) \geq 0$. Since Ψ is increasing, L is increasing on $[1/2, 1]$, so its minimizer is $\tau_h = 1/2$. Together with the preceding $\alpha \leq r$ argument, this proves that no disclosure is optimal whenever $\alpha \leq \alpha^{\mathrm{N}}$. Q.E.D.

Proof of Proposition 3. We first establish the monotonicity of π_l^* within the two partial-disclosure regimes and then combine these results with the policy transitions characterized in Proposition 2.

Step 1: Cutoff partial disclosure. Under Case II(c) of Proposition 2, we have $\frac{d\pi_l^*}{d\alpha} = \frac{1}{(\tau^\dagger)^2}\frac{2r-1}{(2\alpha-1)^2} < 0$ by noticing that $r < \frac{1}{2}$. Hence, π_l^* is decreasing in α.

Step 2: Balanced partial disclosure. Under Cases II(b) and III(b) of Proposition 2, $\pi_l^* = 2 - \frac{1}{\tau^o}$, where τ^o is defined in Proposition 2. Since $2 - \frac{1}{\tau^o}$ is increasing in τ^o, it suffices to show that τ^o is increasing in α.

Define

$$u(\alpha) := \frac{\sqrt{(2\alpha-1)(1-r)^2\sigma - (1-\alpha)^2(\alpha\Delta - \sigma)}}{(1-r)\sqrt{\alpha\sigma}}.$$

Since $\alpha > \frac{1}{2}$ and $\Delta \geq 2\sigma$, we have $\alpha\Delta - \sigma > 0$. Thus,

$$u^2 = \frac{(2\alpha-1)(1-r)^2\sigma - (1-\alpha)^2(\alpha\Delta-\sigma)}{\alpha(1-r)^2\sigma} < \frac{2\alpha-1}{\alpha} < 1,$$

so $0 < u < 1$. Using this definition, the expression for τ^o in Proposition 2 can be rewritten as $\tau^o = \frac{1-\alpha}{1-2\alpha+\frac{2\alpha-1}{u}} = \frac{(1-\alpha)u}{(2\alpha-1)(1-u)}$.

Because balanced partial disclosure requires $\tau^o \leq 1$, and $(2\alpha-1)(1-u) > 0$, we obtain $(1-\alpha)u \leq (2\alpha-1)(1-u)$, or equivalently $\alpha u \leq 2\alpha - 1$. Hence, $0 < u \leq \frac{2\alpha-1}{\alpha}$.

Next, the definition of u implies $(1-r)^2\sigma = \frac{(1-\alpha)^2(\alpha\Delta-\sigma)}{(2\alpha-1)-\alpha u^2}$. Differentiating both sides with respect to α and solving for $u'(\alpha)$ yields

$$u'(\alpha) = \frac{(1-\alpha)^2(\alpha\Delta-\sigma)(2-u^2) + (1-\alpha)(3\alpha\Delta - \Delta - 2\sigma)\big((2\alpha-1) - \alpha u^2\big)}{2\alpha(1-\alpha)^2(\alpha\Delta-\sigma)u}.$$

Differentiating $\tau^o = \frac{(1-\alpha)u}{(2\alpha-1)(1-u)}$ with respect to α gives

$$\frac{d\tau^o}{d\alpha} = \frac{(1-\alpha)(2\alpha-1)u' - u(1-u)}{(2\alpha-1)^2(1-u)^2}.$$

Substituting the expression for $u'(\alpha)$ and simplifying yields

$$\frac{d\tau^o}{d\alpha} = \frac{\mathcal{J}_0(\alpha, u)}{2\alpha u(2\alpha-1)^2(1-u)^2(\alpha\Delta-\sigma)},$$

where

$$\begin{aligned}\mathcal{J}_0(\alpha,u) = &-4\Delta\alpha^3u^2+8\Delta\alpha^3+2\Delta\alpha^2u^3-10\Delta\alpha^2+5\Delta\alpha-\Delta\\ &+2\alpha^2\sigma u^2-4\alpha^2\sigma-2\alpha\sigma u^3+3\alpha\sigma u^2+2\alpha\sigma-\sigma u^2.\end{aligned}$$

Because $\alpha > \frac{1}{2}$, $u > 0$, $u < 1$, and $\alpha\Delta - \sigma > 0$, the denominator is strictly positive. It therefore remains to show that $\mathcal{J}_0(\alpha,u) > 0$.

Let $\eta := \Delta/\sigma \geq 2$. Then $\mathcal{J}_0(\alpha,u) = \sigma\mathcal{J}_1(\alpha,u,\eta)$, where

$$\mathcal{J}_1(\alpha,u,\eta) = 2\alpha(\alpha\eta-1)u^3+(-4\eta\alpha^3+2\alpha^2+3\alpha-1)u^2+(2\alpha-1)(4\eta\alpha^2-3\eta\alpha+\eta-2\alpha).$$

To establish positivity, first consider its derivative with respect to η. Define

$$\mathcal{J}_2(\alpha,u) := \frac{\partial\mathcal{J}_1(\alpha,u,\eta)}{\partial\eta} = 2\alpha^2u^3-4\alpha^3u^2+(2\alpha-1)(4\alpha^2-3\alpha+1).$$

Its derivative with respect to u is $\frac{\partial\mathcal{J}_2}{\partial u} = 2\alpha^2u(3u-4\alpha)$. Since $u \leq \frac{2\alpha-1}{\alpha} < \frac{4\alpha}{3}$, we have $\frac{\partial\mathcal{J}_2}{\partial u} < 0$, so $\mathcal{J}_2(\alpha,u)$ is decreasing in u. Therefore,

$$\mathcal{J}_2(\alpha,u) \geq \mathcal{J}_2\left(\alpha,\frac{2\alpha-1}{\alpha}\right) = \frac{(1-\alpha)(2\alpha-1)(4\alpha^2-5\alpha+2)}{\alpha} > 0,$$

where the last inequality follows because $4\alpha^2-5\alpha+2 > 0$. Hence $\partial\mathcal{J}_1/\partial\eta > 0$. Since $\eta \geq 2$, it follows that $\mathcal{J}_1(\alpha,u,\eta) \geq \mathcal{J}_1(\alpha,u,2)$.

It remains to verify that $\mathcal{J}_1(\alpha,u,2) > 0$. Substituting $\eta = 2$ gives

$$\mathcal{J}_1(\alpha,u,2) = 2\alpha(2\alpha-1)u^3+(-8\alpha^3+2\alpha^2+3\alpha-1)u^2+(2\alpha-1)(8\alpha^2-8\alpha+2).$$

Using $-8\alpha^3+2\alpha^2+3\alpha-1 = (2\alpha-1)(-4\alpha^2-\alpha+1)$ and $8\alpha^2-8\alpha+2 = 2(2\alpha-1)^2$, we obtain $\mathcal{J}_1(\alpha,u,2) = (2\alpha-1)\mathcal{J}_3(\alpha,u)$, where

$$\mathcal{J}_3(\alpha,u) := 2\alpha u^3-(4\alpha^2+\alpha-1)u^2+2(2\alpha-1)^2.$$

Differentiating with respect to u gives $\frac{\partial\mathcal{J}_3}{\partial u} = 2u[3\alpha u-(4\alpha^2+\alpha-1)]$. Since $u \leq \frac{2\alpha-1}{\alpha}$, we have $3\alpha u \leq 3(2\alpha-1)$, and therefore

$$3\alpha u-(4\alpha^2+\alpha-1) \leq 3(2\alpha-1)-(4\alpha^2+\alpha-1) = -(4\alpha^2-5\alpha+2) < 0.$$

Hence $\mathcal{J}_3(\alpha,u)$ is decreasing in u. It follows that

$$\mathcal{J}_3(\alpha,u) \geq \mathcal{J}_3\left(\alpha,\frac{2\alpha-1}{\alpha}\right) = \frac{(1-\alpha)(2\alpha-1)^3}{\alpha^2} > 0.$$

Thus $\mathcal{J}_1(\alpha, u, 2) > 0$. Since $\mathcal{J}_1$ is increasing in η and $\eta \geq 2$, we have $\mathcal{J}_1(\alpha, u, \eta) > 0$, which implies $\mathcal{J}_0(\alpha, u) > 0$. Consequently, $\frac{d\tau^o}{d\alpha} > 0$. Hence τ^o, and therefore $2 - \frac{1}{\tau^o}$, is increasing in α throughout the balanced partial-disclosure region.

Step 3: Global monotonicity. Under full and no disclosure, π_l^* is constant at 1 and 0, respectively. For $r < 1/2$, if $\alpha^{\mathrm{C}} \leq 1 - r$, the proof of Proposition 2 gives $\alpha^{\mathrm{F}} \leq 1 - r$, and full disclosure is optimal for every $\alpha \in (1/2, 1)$. Thus, the weak monotonicity claims in part (i) hold. If $\alpha^{\mathrm{C}} > 1 - r$, then $1 - r < \alpha^{\mathrm{C}} < \alpha^{\mathrm{F}}$, so Steps 1 and 2 imply that π_l^* strictly decreases in the cutoff partial-disclosure region and strictly increases in the balanced partial-disclosure region, while remaining constant in the full-disclosure regions. Since $G_{\mathrm{C}}(1-r) = \sigma[r^2 + (1-r)^2] - \Delta r(1-r)$, this non-monotonicity occurs precisely when $\Delta r(1-r) > \sigma[r^2 + (1-r)^2]$.

For $r \geq 1/2$, Step 2 and the constant policies imply that π_l^* is weakly increasing in α. Since adjacent policies coincide at their boundaries, these arguments establish both parts. Q.E.D.

Proof of Proposition 4. For part (i), define $\sigma_0 = \alpha(1-\alpha)\Delta/[(1-\alpha) + (4\alpha^2 - 1)(1-r)^2]$, $\sigma_1 = \alpha(1-\alpha)\Delta/[\alpha^2 + 2(1-\alpha)r - r^2]$, and $\sigma_2 = \alpha^2(1-\alpha)\Delta/[\alpha(1-\alpha) + (2\alpha - 1)(1-r)^2]$. Equation (A.2) gives $\Psi(1/2) = (1-\alpha)\Delta(1-\sigma/\sigma_0)$, $\Psi(\tau^\dagger) = (1-\alpha)\Delta(1-\sigma/\sigma_1)$, and $\Psi(1) = (1-\alpha)\Delta(1-\sigma/\sigma_2)$. Since Ψ is strictly increasing in (A.1), its values at $\tau^\dagger \leq 1$ imply $\sigma_1 \leq \sigma_2$ in Case II; its values at $1/2 < 1$ imply $\sigma_0 < \sigma_2$ in Case III. At $\sigma = \sigma_2$, $\Psi(1) = 0$, so $\tau_h = 1$ implements full disclosure. These endpoint signs, together with Proposition 2, yield the stated policy intervals.

For part (ii), full disclosure, no disclosure, and cutoff partial disclosure set π_l^* equal to 1, 0, and $2 - 1/\tau^\dagger$, respectively, all independent of σ. Under balanced disclosure, let $\mu^o = \mu(\tau^o) = 1 - \alpha + (2\alpha - 1)\tau^o$. Define $T(\mu) = (1-r)^2(2\mu + \alpha - 1)/\mu^2 - [(1-r)^2 - (1-\alpha)]/\alpha$. The first-order condition in (A.1) becomes $\sigma T(\mu^o) = (1-\alpha)\Delta$. Since $T'(\mu) = -2(1-r)^2(\mu + \alpha - 1)/\mu^3 < 0$, implicit differentiation with Δ fixed gives $d\mu^o/d\sigma = -T(\mu^o)/[\sigma T'(\mu^o)] > 0$. Hence $d\pi_l^*/d\sigma = [(2\alpha - 1)(\tau^o)^2]^{-1} d\mu^o/d\sigma > 0$.

For part (iii), substituting the quantities from Lemma 1 into the supplier's profit and applying the Bayes-plausible message probabilities gives its expected payoff under balanced

disclosure:

$$\begin{aligned}\Pi_S^{\mathrm{B}} = w\left(1-\frac{1}{2\tau^o}\right)r\left[\theta_L-\frac{\sigma(\alpha-r)}{\alpha}\right] \\ +\frac{w}{2\tau^o}\left\{r\theta_L+\Delta(\mu^o-1+r)+\sigma\left[\frac{(1-r)^2}{\mu^o}-(1-r)\right]\right\}.\end{aligned} \tag{A.3}$$

Differentiating (A.3), including the changes in both the posterior and its probability, and then substituting $\sigma T(\mu^o) = (1-\alpha)\Delta$ yields $d\Pi_S^{\mathrm{B}}/d\sigma = wK(\mu^o)$, where

$$\begin{aligned}K(\mu) = &-\frac{(2\mu-1)r(\alpha-r)}{2\alpha(\mu+\alpha-1)}+\frac{2\alpha-1}{2(\mu+\alpha-1)}\left[\frac{(1-r)^2}{\mu}-(1-r)\right] \\ &+\frac{(2\alpha-1)\mu^3T(\mu)}{4(1-r)^2(\mu+\alpha-1)^3}\left[\frac{(2\alpha-r-1)T(\mu)}{1-\alpha}+\frac{2r(1-\alpha)}{\alpha}\right].\end{aligned} \tag{A.4}$$

We show that K in (A.4) is positive throughout the relevant posterior interval. Since T is decreasing and $T(\alpha) = [\alpha(1-\alpha)+(2\alpha-1)(1-r)^2]/\alpha^2 > 0$, we have $T(\mu) > 0$ for $\max\{1/2, 1-r\} \le \mu \le \alpha$. Direct differentiation gives

$$K'(\mu) = -\frac{(2\alpha-1)T(\mu)\left[3(2\alpha-1)\mu^2(\alpha-r)^2+\alpha(1-r)^2(\mu+\alpha-1)^2\right]}{4(1-\alpha)\alpha(1-r)(\mu+\alpha-1)^4} < 0. \tag{A.5}$$

Moreover,

$$\begin{aligned}&4(1-\alpha)\alpha(2\alpha-1)(1-r)K(\alpha) \\ &\quad=(2\alpha-1)\left[(1-r)^2+(1-\alpha)(1-r)-(1-\alpha)\right]^2+(1-\alpha)^2(\alpha-r)^2 \\ &\qquad+2(1-\alpha)^2(2\alpha-1)(\alpha-r)(\alpha+r-1)+2\alpha(1-\alpha)^3(2\alpha-1)>0.\end{aligned} \tag{A.6}$$

Thus (A.5)–(A.6) imply $K(\mu^o) \ge K(\alpha) > 0$, and the supplier's payoff is strictly increasing throughout the balanced interval.

For the retailer, expected demand is independent of σ, and τ^o minimizes the expected shortage cost L in (A.1). The envelope theorem therefore gives

$$\frac{d\Pi_R^{\mathrm{Opt}}}{d\sigma} = -\frac{p-w}{2(\mu^o+\alpha-1)}\left\{\frac{(2\mu^o-1)[(1-r)^2-(1-\alpha)]}{\alpha}+\frac{(2\alpha-1)(1-r)^2}{\mu^o}\right\}. \tag{A.7}$$

To establish the sign in (A.7), rewrite the first-order condition as $(1-r)^2[\alpha(2\mu^o+\alpha-1)-(\mu^o)^2] = (1-\alpha)(\mu^o)^2(\alpha\Delta/\sigma-1)$. Its bracket is positive because $\Delta/\sigma \ge 2$. Moreover, $(2\alpha-1)\mu^o(\alpha+\mu^o)-[\alpha(2\mu^o+\alpha-1)-(\mu^o)^2] = \alpha(2\mu^o-1)(\alpha+\mu^o-1) > 0$. Combining these relations yields $(1-r)^2 > (1-\alpha)\mu^o/(\alpha+\mu^o)$. The expression in braces in (A.7) is strictly increasing in $(1-r)^2$; at this lower bound it equals $2(1-\alpha)(\alpha-\mu^o)/(\alpha+\mu^o) \ge 0$. Thus $d\Pi_R^{\mathrm{Opt}}/d\sigma < 0$.

For the supply chain, define

$$\hat{p}(\sigma) := w + \frac{w(p-w)K(\mu^o)}{-d\Pi_R^{\text{Opt}}/d\sigma} > w. \tag{A.8}$$

Since the supply chain payoff equals $\Pi_R^{\text{Opt}} + \Pi_S^{\text{B}}$, its derivative is

$$\frac{d}{d\sigma}\left(\Pi_R^{\text{Opt}} + \Pi_S^{\text{B}}\right) = \frac{\hat{p}(\sigma) - p}{p - w}\left(-\frac{d\Pi_R^{\text{Opt}}}{d\sigma}\right),$$

which is positive if and only if $p < \hat{p}(\sigma)$. Expected sales and production cost equal $\Pi_R^{\text{Opt}}/(p-w)$ and $w\Pi_R^{\text{Opt}}/(p-w) - \Pi_S^{\text{B}}$, respectively. Thus, $\hat{p}(\sigma)$ is precisely the production cost saved per unit reduction in expected sales. This cutoff is independent of p, since $p-w$ only scales the retailer's objective and cancels from the ratio in (A.8). Finally, $\hat{p}(\sigma) > w$ ensures a nonempty interval $w < p < \hat{p}(\sigma)$ in which the supply chain benefits. Q.E.D.

LEMMA A.1. *Suppose that $0 < \Delta < 2\sigma$. For any belief μ, the supplier's optimal production quantity Q_μ^* is given by*

$$Q_\mu^* = \begin{cases} \theta_L + \sigma\left(\dfrac{2r}{1-\mu} - 1\right), & \mu \le \mu_1, \\ \theta_L - \sigma + 2\sigma r + \Delta\mu, & \mu_1 < \mu < \mu_2, \\ \theta_H + \sigma\left(1 - \dfrac{2(1-r)}{\mu}\right), & \mu \ge \mu_2, \end{cases}$$

where $\mu_1 := 1 - 2r\sigma/\Delta$ and $\mu_2 := 2(1-r)\sigma/\Delta$.

Proof of Lemma A.1. Fix a posterior belief $\mu \in [0,1]$, and let $F_D(\cdot;\mu)$ denote the conditional c.d.f. of demand. Since $0 < \Delta < 2\sigma$, we have $\theta_L - \sigma < \theta_H - \sigma < \theta_L + \sigma < \theta_H + \sigma$. Thus,

$$F_D(x;\mu) = \begin{cases} 0, & x < \theta_L - \sigma, \\ (1-\mu)\dfrac{x - (\theta_L - \sigma)}{2\sigma}, & \theta_L - \sigma \le x < \theta_H - \sigma, \\ \dfrac{x - (\theta_L - \sigma) - \mu\Delta}{2\sigma}, & \theta_H - \sigma \le x \le \theta_L + \sigma, \\ 1 - \mu + \mu\dfrac{x - (\theta_H - \sigma)}{2\sigma}, & \theta_L + \sigma < x \le \theta_H + \sigma, \\ 1, & x > \theta_H + \sigma. \end{cases}$$

The supplier's optimal production quantity is the r-quantile of this distribution. At the two boundaries of the overlap interval,

$$F_D(\theta_H - \sigma; \mu) = \frac{(1-\mu)\Delta}{2\sigma}, \qquad F_D(\theta_L + \sigma; \mu) = 1 - \frac{\mu\Delta}{2\sigma}.$$

Hence, the r-quantile lies in the low-, overlap-, and high-production intervals according as $\mu \le \mu_1$, $\mu_1 < \mu < \mu_2$, and $\mu \ge \mu_2$, respectively. Solving $F_D(Q;\mu) = r$ on the corresponding branches gives the three expressions for Q^*_μ stated in the lemma. Finally, $\mu_2 - \mu_1 = (2\sigma - \Delta)/\Delta > 0$, so the two breakpoints are correctly ordered. Q.E.D.

Proof of Proposition 5. By Corollary A.1, the retailer minimizes expected shortage cost. Under overlapping demand, the cost associated with posterior μ is

$$V_{\text{ov}}(\mu) = \begin{cases} \Delta\mu + \dfrac{\sigma\left[(1-r)^2 - \mu\right]}{1-\mu}, & \mu \le \mu_1, \\ \sigma(1-r)^2 + \dfrac{\Delta^2\mu(1-\mu)}{4\sigma}, & \mu_1 < \mu < \mu_2, \\ \dfrac{\sigma(1-r)^2}{\mu}, & \mu \ge \mu_2. \end{cases} \tag{A.9}$$

The optimal expected cost is the value at $\tau = 1/2$ of the lower convex envelope of $V_{\text{ov}}(\mu(\tau))$. We again first identify all possible contact points of this envelope, then derive the critical ratios at which its supporting chord changes, and finally translate each supporting chord into a disclosure policy.

We first consider $\alpha = 1$. Full disclosure induces posteriors $\mu = 0$ and $\mu = 1$, and $V_{\text{ov}}(0) = V_{\text{ov}}(1) = \sigma(1-r)^2$. We show that this is the global minimum of V_{ov}. On the first branch, $V_{\text{ov}}(\mu) - \sigma(1-r)^2 = \mu[\Delta(1-\mu) - \sigma r(2-r)]/(1-\mu) \ge 0$, because $\mu \le 1 - 2r\sigma/\Delta$ implies $\Delta(1-\mu) \ge 2r\sigma \ge \sigma r(2-r)$. On the second branch, $V_{\text{ov}}(\mu) - \sigma(1-r)^2 = \Delta^2\mu(1-\mu)/(4\sigma) \ge 0$, and on the third branch, $V_{\text{ov}}(\mu) - \sigma(1-r)^2 = \sigma(1-r)^2(1/\mu - 1) \ge 0$. Hence $V_{\text{ov}}(\mu) \ge \sigma(1-r)^2$ for all $\mu \in [0,1]$, and full disclosure attains this lower bound. Therefore, full disclosure is optimal when $\alpha = 1$.

Henceforth, suppose $\alpha \in (1/2, 1)$. As in the baseline proof, parameterize feasible posteriors by $\mu(\tau) = 1 - \alpha + (2\alpha - 1)\tau$. The two production breakpoints in the posterior index are $\tau_1 := \frac{\alpha - 2r\sigma/\Delta}{2\alpha - 1}$ and $\tau_2 := \frac{\alpha - 1 + 2(1-r)\sigma/\Delta}{2\alpha - 1}$. They correspond, respectively, to the supplier entering and leaving the overlap production interval. Moreover, $\tau_2 - \tau_1 = (2\sigma/\Delta - 1)/(2\alpha - 1) > 0$.

The three second derivatives of $V_{\text{ov}}(\mu(\tau))$ are $-2\sigma r(2-r)(2\alpha-1)^2/[1-\mu(\tau)]^3 < 0$, $-\Delta^2(2\alpha-1)^2/(2\sigma) < 0$, and $2\sigma(1-r)^2(2\alpha-1)^2/\mu(\tau)^3 > 0$, respectively. Thus the first two branches are strictly concave and the third is strictly convex. An interior point of either concave branch cannot be an isolated contact point of the lower convex envelope. Therefore, a nondegenerate supporting chord that spans the prior can use only 0 or the breakpoint τ_1 as its left contact. Its right contact is either the endpoint 1 or a point on the third, strictly convex branch; a line issued from 0 can have at most one interior tangency point on that branch.

We next rule out any chord from τ_1 to the third branch except possibly the endpoint 1. Suppose $0 < \tau_1 < 1/2 < \widehat{\tau} < 1$ with $\widehat{\tau} \geq \tau_2$, and let $\widehat{\mu} = \mu(\widehat{\tau})$. If $\{\tau_1, \widehat{\tau}\}$ formed part of the lower convex envelope, convexity would require the slope from 0 to τ_1 not to exceed the derivative at $\widehat{\tau}$. These slopes are $-\Delta(2\alpha-1)(2-r-2\alpha)/(2\alpha)$ and $-\sigma(1-r)^2(2\alpha-1)/\widehat{\mu}^2$, respectively.

Because $\widehat{\mu} \leq \alpha$, it is sufficient to show that $\frac{\Delta}{2\sigma}\alpha(2-r-2\alpha) - (1-r)^2 < 0$. The left-hand side is strictly decreasing in α, since its derivative is $\frac{\Delta}{2\sigma}(2-r-4\alpha) < 0$ for $\alpha > 1/2$. If $r \geq 1/2$, the condition $\tau_1 > 0$ implies $\alpha > 2r\sigma/\Delta$. Evaluating the expression at this lower bound gives a value strictly below $-(2r-1)^2 \leq 0$, and hence it is negative at α. If $r < 1/2$, $\widehat{\tau} \geq \tau_2$ and $\widehat{\tau} < 1$ imply $\tau_2 < 1$, so $\alpha > 2(1-r)\sigma/\Delta$. Evaluating at this lower bound gives $(1-r)[1-4(1-r)\sigma/\Delta] < 0$, where the inequality follows from $r < 1/2$ and $\Delta < 2\sigma$. Thus, in either case, the slope from 0 to τ_1 is strictly greater than the supporting slope of the third branch, contradicting the nondecreasing-slope property of a convex function.

Therefore, no chord connecting τ_1 to any $\widehat{\tau} \in [\tau_2, 1)$ that spans the prior can lie on the lower convex envelope. The endpoint chord $\{\tau_1, 1\}$, which may support the downward-cutoff policy, remains a candidate. It follows that the envelope at the prior can be generated only by the endpoint chord $\{0, 1\}$, one of the cutoff chords $\{0, \tau_1\}$ and $\{\tau_1, 1\}$, a supporting line from 0 to the third branch, or the degenerate posterior $1/2$.

We now determine when τ_1 itself is retained. If $\tau_2 \geq 1$, equivalently $r \leq 1 - \alpha\Delta/(2\sigma)$, the right endpoint lies on the second branch. If $\tau_2 < 1$, it lies on the third branch. Direct

calculation gives

$$\operatorname{sgn}\left\{\begin{array}{l}\dfrac{V_{\mathrm{ov}}(\mu(1))-V_{\mathrm{ov}}(\mu(\tau_1))}{1-\tau_1}\\ \quad -\dfrac{V_{\mathrm{ov}}(\mu(\tau_1))-V_{\mathrm{ov}}(\mu(0))}{\tau_1}\end{array}\right\}=\begin{cases}\operatorname{sgn}\{2\sigma(1-\alpha)(2-r)-\Delta\alpha^2\}, & \tau_2\geq 1,\\ \operatorname{sgn}\{2\sigma(1-\alpha)-\Delta[2(1-\alpha)^2+r(2\alpha-1)]\}, & \tau_2<1.\end{cases} \tag{A.10}$$

The lower convex envelope retains τ_1 precisely when the expression on the left of (A.10) is positive. Solving the two equality conditions gives $r^{\mathrm{C}}_{\mathrm{ov},12}(\alpha):=2-\Delta\alpha^2/[2\sigma(1-\alpha)]$ and $r^{\mathrm{C}}_{\mathrm{ov},123}(\alpha):=2(1-\alpha)[\sigma-\Delta(1-\alpha)]/[\Delta(2\alpha-1)]$. We therefore set $r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)=r^{\mathrm{C}}_{\mathrm{ov},12}(\alpha)$ when $\Delta\alpha(2\alpha-1)\geq 2\sigma(1-\alpha)$, and $r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)=r^{\mathrm{C}}_{\mathrm{ov},123}(\alpha)$ otherwise. To verify the branch selection, note that $\tau_2\geq 1$ if and only if $r\leq 1-\alpha\Delta/(2\sigma)$. The differences between the two candidate roots and this branch boundary have the same sign: after removing positive denominators, both signs are determined by $2\sigma(1-\alpha)-\Delta\alpha(2\alpha-1)$. Thus the definition of $r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)$ selects the root that lies in the branch from which it was derived. Consequently, τ_1 is strictly retained when $r<r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)$; at equality, the cutoff chord and the endpoint chord have the same cost.

We next consider the case in which τ_1 is not retained. If $r\geq 1-(1-\alpha)\Delta/(2\sigma)$, then $1-\alpha$ lies on the third branch, so the entire feasible posterior interval $[1-\alpha,\alpha]$ lies on this strictly convex branch. Jensen's inequality then implies that no disclosure is optimal. Hence, it remains to consider $r<1-(1-\alpha)\Delta/(2\sigma)$, for which $1-\alpha$ lies on either the first or the second branch.

The remaining informative candidate is a line from $(0,V_{\mathrm{ov}}(1-\alpha))$ tangent to the third branch. Let τ^o_{ov} denote its candidate contact point and write $\mu^o_{\mathrm{ov}}:=\mu(\tau^o_{\mathrm{ov}})$. The tangency condition reduces to

$$V_{\mathrm{ov}}(1-\alpha)(\mu^o_{\mathrm{ov}})^2-2\sigma(1-r)^2\mu^o_{\mathrm{ov}}+\sigma(1-r)^2(1-\alpha)=0, \tag{A.11}$$

where μ^o_{ov} is the larger root and is retained only when it lies on the third branch, $\mu^o_{\mathrm{ov}}\geq 2(1-r)\sigma/\Delta$. The discriminant of (A.11) is $4\sigma(1-r)^2\{\sigma(1-r)^2-(1-\alpha)V_{\mathrm{ov}}(1-\alpha)\}$. Whenever the tangent exists, evaluating the left-hand side at $\mu=1-\alpha$ shows that the smaller root is weakly below $1-\alpha$. Therefore, the signs of the quadratic at $\mu=1/2$ and $\mu=\alpha$ locate the larger root within the feasible interval.

Accordingly, define $H_{\mathrm{F}}(r) := V_{\mathrm{ov}}(1-\alpha)\alpha^2 - \sigma(1-r)^2(3\alpha-1)$ and $H_{\mathrm{N}}(r) := V_{\mathrm{ov}}(1-\alpha) - 4\sigma(1-r)^2\alpha$. Substituting $\mu^o_{\mathrm{ov}} = \alpha$ and $\mu^o_{\mathrm{ov}} = 1/2$ into (A.11) gives $H_{\mathrm{F}}(r) = 0$ and $H_{\mathrm{N}}(r) = 0$, respectively. Hence $\tau^o_{\mathrm{ov}} \geq 1$ if and only if $H_{\mathrm{F}}(r) \leq 0$, $1/2 < \tau^o_{\mathrm{ov}} < 1$ if and only if $H_{\mathrm{F}}(r) > 0$ and $H_{\mathrm{N}}(r) < 0$, and $\tau^o_{\mathrm{ov}} \leq 1/2$ if and only if $H_{\mathrm{N}}(r) \geq 0$.

Suppose first that $r \leq \alpha\Delta/(2\sigma)$, so that $1-\alpha$ lies on the first branch. Substitution gives $H_{\mathrm{F}}(r) = \alpha(1-\alpha)(\alpha\Delta - \sigma) - \sigma(2\alpha-1)(1-r)^2$ and $\alpha H_{\mathrm{N}}(r) = (1-\alpha)(\alpha\Delta-\sigma) - \sigma(4\alpha^2-1)(1-r)^2$. If $\alpha\Delta > \sigma$, their respective zeros are $r^{\mathrm{F}}_{\mathrm{ov},123}(\alpha) := 1 - \sqrt{\alpha(1-\alpha)(\alpha\Delta-\sigma)/[\sigma(2\alpha-1)]}$ and $r^{\mathrm{N}}_{\mathrm{ov},123}(\alpha) := 1 - \sqrt{(1-\alpha)(\alpha\Delta-\sigma)/[\sigma(4\alpha^2-1)]}$. If $\alpha\Delta \leq \sigma$, both comparisons are strictly negative throughout this branch, so neither boundary equation has a branch-consistent root.

Next suppose that $\alpha\Delta/(2\sigma) < r < 1 - (1-\alpha)\Delta/(2\sigma)$, so that $1-\alpha$ lies on the second branch. In this case, $H_{\mathrm{F}}(r) = \Delta^2\alpha^3(1-\alpha)/(4\sigma) - \sigma(3\alpha - 1 - \alpha^2)(1-r)^2$ and $H_{\mathrm{N}}(r) = \Delta^2\alpha(1-\alpha)/(4\sigma) - \sigma(4\alpha-1)(1-r)^2$. Their respective zeros are $r^{\mathrm{F}}_{\mathrm{ov},23}(\alpha) := 1 - \frac{\Delta}{2\sigma}\sqrt{\alpha^3(1-\alpha)/(3\alpha-1-\alpha^2)}$ and $r^{\mathrm{N}}_{\mathrm{ov},23}(\alpha) := 1 - \frac{\Delta}{2\sigma}\sqrt{\alpha(1-\alpha)/(4\alpha-1)}$.

On the first branch, the derivatives of H_{F} and H_{N} are $2\sigma(2\alpha-1)(1-r) > 0$ and $2\sigma(4\alpha^2-1)(1-r)/\alpha > 0$, respectively. On the second branch, they are $2\sigma(3\alpha-1-\alpha^2)(1-r) > 0$ and $2\sigma(4\alpha-1)(1-r) > 0$. Because $V_{\mathrm{ov}}(1-\alpha)$ is continuous at $r = \alpha\Delta/(2\sigma)$, both comparisons are continuous and strictly increasing across the two branches. Hence, for $j \in \{\mathrm{F}, \mathrm{N}\}$, the branch-consistent threshold is $r^j_{\mathrm{ov},23}(\alpha)$ when $\alpha\Delta \leq \sigma$, and $\min\{r^j_{\mathrm{ov},23}(\alpha), r^j_{\mathrm{ov},123}(\alpha)\}$ otherwise.

Finally, $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha) < r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$. At $r = r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$, $H_{\mathrm{N}}(r) = 0$ implies $V_{\mathrm{ov}}(1-\alpha) = 4\sigma(1-r)^2\alpha$. Substituting this equality into H_{F} gives $H_{\mathrm{F}}(r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)) = \sigma(1-r)^2[4\alpha^3 - 3\alpha + 1] = \sigma(1-r)^2(2\alpha-1)^2(\alpha+1) > 0$. Since H_{F} is strictly increasing and its unique relevant zero is $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha)$, it follows that $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha) < r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$. Moreover, direct substitution into the one-sided supporting condition at τ_2 shows that, whenever $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha) < r < r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$, the larger root of (A.11) satisfies $\mu^o_{\mathrm{ov}} \geq 2(1-r)\sigma/\Delta$. Thus the interior tangency selected by these thresholds lies on the third branch, as required; equality corresponds only to a boundary contact and creates no additional policy region.

We now prove the proposition case by case.

Case I: $r < (1-\alpha)\Delta/(2\sigma)$. This condition is equivalent to $\tau_1 > 1$. The feasible interval therefore lies entirely on the first strictly concave branch of V_{ov}, whose lower convex envelope is the chord joining 0 and 1. Hence full disclosure is optimal.

Case II: $(1-\alpha)\Delta/(2\sigma) \le r < \Delta/(4\sigma)$. Here $1/2 < \tau_1 \le 1$. If $r < r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)$, equation (A.10) shows that τ_1 is retained, so the lower convex envelope at the prior is the chord joining 0 and τ_1. The Bayes-plausibility calculation in the proof of Proposition 2 therefore gives the upward-cutoff policy in part (a). At the lower boundary, $\tau_1 = 1$, so it coincides with full disclosure.

If $r \ge r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)$, the breakpoint is not retained. A balanced tangent cannot replace it in Case II. Indeed, $r < \Delta/(4\sigma) < \alpha\Delta/(2\sigma)$, so $V_{\mathrm{ov}}(1-\alpha)$ uses the first branch. The full-disclosure boundary expression obtained from (A.11) is $\alpha(1-\alpha)(\alpha\Delta-\sigma) - \sigma(2\alpha-1)(1-r)^2$. It is negative immediately if $\alpha\Delta \le \sigma$. If $\alpha\Delta > \sigma$, then $(\alpha\Delta-\sigma)/[\sigma(2\alpha-1)] < 1$, so its first term after division by $\sigma(2\alpha-1)$ is smaller than $\alpha(1-\alpha) \le 1/4$, whereas $(1-r)^2 > (1-\Delta/(4\sigma))^2 > 1/4$. Thus $\tau^o_{\mathrm{ov}} > 1$, and the endpoint chord joining 0 and 1 is the lower convex envelope at the prior. This proves full disclosure in part (b).

Case III: $r \ge \Delta/(4\sigma)$. Here $\tau_1 \le 1/2$. The downward-cutoff chord is feasible at the prior while $\tau_1 > 0$, or equivalently $r < \alpha\Delta/(2\sigma)$, and it belongs to the lower convex envelope exactly when $r < r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)$. Define $r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha) := \min\{\alpha\Delta/(2\sigma), r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)\}$. The two conditions then combine to $r < r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha)$, proving part (a).

The remaining thresholds satisfy $r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha) < r^{\mathrm{F}}_{\mathrm{ov}}(\alpha) < r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$. Indeed, because $r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha) \le \alpha\Delta/(2\sigma)$, the first-branch expression for H_{F} applies at this cutoff. Substituting $r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha) = \min\{\alpha\Delta/(2\sigma), r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)\}$, together with the two branch formulas for $r^{\mathrm{C}}_{\mathrm{ov}}(\alpha)$, gives $H_{\mathrm{F}}(r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha)) < 0$ in every admissible case. Since H_{F} is strictly increasing and $H_{\mathrm{F}}(r^{\mathrm{F}}_{\mathrm{ov}}(\alpha)) = 0$, it follows that $r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha) < r^{\mathrm{F}}_{\mathrm{ov}}(\alpha)$. Combining this with $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha) < r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$, established above, yields the stated ordering.

Therefore, for $r \ge r^{\mathrm{DC}}_{\mathrm{ov}}(\alpha)$, the endpoint chord is optimal through $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha)$, an interior tangent $\{0, \tau^o_{\mathrm{ov}}\}$ is optimal between $r^{\mathrm{F}}_{\mathrm{ov}}(\alpha)$ and $r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$, and the degenerate posterior is optimal thereafter. This gives parts (b)–(d). The implementation of $\{0, \tau^o_{\mathrm{ov}}\}$ and of no disclosure is the same as in Proposition 2. At $r = r^{\mathrm{F}}_{\mathrm{ov}}(\alpha)$, the tangency point is 1, so balanced partial disclosure coincides with full disclosure; at $r = r^{\mathrm{N}}_{\mathrm{ov}}(\alpha)$, it is $1/2$, so balanced partial disclosure coincides with no disclosure. Q.E.D.

B Supplemental Results and Proofs

LEMMA B.1. *Under balanced partial disclosure, the high-message posterior μ_h is increasing in forecast capability α.*

Proof of Lemma B.1. Under balanced partial disclosure, $\mu_h = 1-\alpha+(2\alpha-1)\tau^o$. Step 2 of the proof of Proposition 3 establishes that $d\tau^o/d\alpha > 0$. Since $\tau^o \geq 1/2$ and $\alpha > 1/2$, $\frac{d\mu_h}{d\alpha} = 2\tau^o - 1 + (2\alpha-1)\frac{d\tau^o}{d\alpha} > 0$. Q.E.D.

COROLLARY B.1. *For $\Delta \geq 2\sigma$, consider intervals on which no disclosure, cutoff partial disclosure, or full disclosure is optimal.*

(i) If $1-\sqrt{\alpha^2+(1-\alpha)^2} < r < 1-\alpha$, the retailer's optimal profit increases in σ, and the supply chain's expected profit increases if and only if $p > \underline{p}^{\mathrm{F}}$, where $\underline{p}^{\mathrm{F}} = 2c[r-\alpha(1-\alpha)]/[r(2-r)-2\alpha(1-\alpha)]$.

(ii) For all other values of r, the retailer's optimal payoff and the supply chain's expected profit are decreasing in σ.

Proof of Corollary B.1. For part (i), full disclosure is optimal by Proposition 2. Lemmas 1 and 2 give $d\Pi_R^{\mathrm{Opt}}/d\sigma = (p-w)[r(2-r)-2\alpha(1-\alpha)]/[2\alpha(1-\alpha)] > 0$. Adding the supplier's payoff derivative gives the supply chain's derivative $(p-\underline{p}^{\mathrm{F}})[r(2-r)-2\alpha(1-\alpha)]/[2\alpha(1-\alpha)]$. The fraction is positive on the stated interval, proving part (i).

For part (ii), if $r \leq 1-\sqrt{\alpha^2+(1-\alpha)^2}$, full disclosure remains optimal. The retailer's derivative above is negative, whereas the supplier's derivative is $wr[r-2\alpha(1-\alpha)]/[2\alpha(1-\alpha)] < 0$. Hence the supply chain's payoff strictly decreases.

Under no disclosure, $r \geq 1/2$, and the retailer's and supplier's derivatives are $-2(p-w)(1-r)^2 < 0$ and $c(1-2r) \leq 0$, respectively. Under cutoff partial disclosure, $1-\alpha \leq r < 1/2$, and the retailer's derivative is $-(p-w)\{(1-2r)(2-r)/(2\alpha)+(\alpha+r-1)(1-r)/(\alpha-r)\} < 0$, while the supplier's derivative is $wr(2r-1)/(2\alpha) < 0$. Thus both payoffs in part (ii) decrease under these policies.

For optimal full disclosure with $r \geq 1-\alpha$, necessarily $r < \alpha$. Optimality gives $G_{\mathrm{F}}(\alpha) \geq 0$; by (A.2) and $\Delta/\sigma \geq 2$, this implies $(1-r)^2 \geq \alpha(1-\alpha)$. The retailer's derivative is therefore $(p-w)[1-\alpha-2(1-r)^2]/(2\alpha) < 0$. The supplier's derivative is $w[1-\alpha-2r(1-r)]/(2\alpha) < 0$, since $r(1-r) \geq \alpha(1-\alpha)$. Their sum is strictly negative, completing part (ii). Q.E.D.