Two-sided Benefits of Price Transparency in Smallholder Supply Chains

Yuan Shi
MIT Operations Research Center, yuansh@mit.edu

Joann F. de Zegher
MIT Sloan School of Management, Operations Management, jfz@mit.edu

Irene Y. Lo
Stanford School of Engineering, Management Science & Engineering, ilo@stanford.edu

Information platforms have emerged in the developing world to improve price transparency and welfare for smallholder suppliers. Meanwhile, sustaining welfare improvement often requires such platforms to benefit both suppliers and buyers. This paper studies the impact of price transparency on market price and welfare in smallholder supply chains, and identifies conditions and driving forces for creating benefits to both suppliers and buyers. Motivated by granular data from smallholder supply chains, we develop a new Hotelling model of price search, where price-setting buyers face the operational challenges of demand asymmetry and costly underage or overage amid uncertain supply. We find that high overage costs, combined with high demand asymmetry that dominates random supply variations, give rise to two-sided benefits, driven by price competition benefiting the suppliers and demand signaling benefiting the buyers under increased transparency. Moreover, achieving two-sided benefit requires implementing a well-chosen level of price transparency, and, in some cases, creating a low uncertainty environment for buyers. These results help close the gap between the empirical literature and the theoretical economic literature on this topic, and offer possible explanations for the variation in empirical findings. We provide managerial recommendations for information platform designers, including for our partnering platform, on identifying the target markets and whether to implement full or partial transparency.

Key words: smallholder supply chains, information platforms, newsvendor, price search, collusion

---

1. Introduction

Smallholder farms up to 2 hectares in size produce one-third of the world's food, and those up to 20 hectares produce close to 60% of global food calories consumed (Ritchie 2021). These smallholder farmers sell through supply chains that are opaque, with limited access to market information (*e.g.*,Sodhi and Tang 2014, Chen and Tang 2015). As a result, price-searching suppliers (smallholder farmers) cannot get the best available price, and price-setting buyers cannot utilize pricing as an efficient signaling mechanism to source the supply they need when they need it.

To mediate such inefficiencies, mobile-based platforms have emerged to improve price transparency in smallholder supply chains, by providing suppliers with access to prices outside their informal network. Our work is inspired by such a platform, *PemPem*, used by smallholder suppliers and buyers in the palm oil supply chain in rural Indonesia. *PemPem* crowdsources price information for smallholder palm fruit farmers and middlemen, who had relied on limited personal connections to trade with buyers prior to the introduction of the platform. Other digital platforms that improve price transparency in smallholder supply chains in the developing world include *Loop* and *Reuters Market Light (RML)* in India and *Kudu* in Uganda Newman et al. (2018).

Most existing works in operations management studying information provision for smallholders have focused on poverty alleviation for the economically vulnerable suppliers, largely ignoring the welfare impact on buyers (see, e.g. Sodhi and Tang 2014, Chen and Tang 2015, Liao et al. 2019, Zhou et al. 2021). However, to encourage widespread adoption of information platforms and ensure sustained welfare benefits, it is often crucial that both sides of the market benefit, such as when a platform relies on crowd-sourcing to collect information. In the case of *PemPem*, while both suppliers and buyers are encouraged to share price information on a daily basis for others to use, the price-setting buyers provide >90% of daily price information. Price-seeking suppliers are reluctant to share information due to competition among themselves, which is a phenomenon commonly observed in literature (see, e.g. Chen et al. 2015, Xiao et al. 2020). However, if too much transparency leads to elevated price competition among the buyers, they may also become reluctant to share the information voluntarily. On *PemPem*, while some buyers are willing to share information, others are not because they believe that increased transparency hurts their business due to increased competition. For such a platform, a sustainable information provision strategy must benefit both sides: it should deliver enough value to sellers while encouraging buyers, its key provider of price information, to continue sharing data.

Motivated by this objective, this paper studies whether and how two-sided benefits can arise from increased price transparency in smallholder supply chains. We develop a novel Hotelling duopoly model where buyers compete for supply from sellers through pricing and face supply uncertainty and contractual obligations. In the absence of a platform, each buyer has informal relationships with a set of loyal sellers. While smallholder farmers do not face binding obligations to sell to the specific buyer, their selling options are restricted due to limited information and market access. Introduction of the platform creates price transparency and enables price search by the sellers, who choose among the competing buyers based on price and transportation costs. Figure 1 demonstrates such price searching behavior in Indonesia's palm oil supply chain using data from *PemPem*: palm fruit farmers using the platform trade more with buyers within geographical proximity and those who bid higher prices.

We introduce in our model two key operational issues faced by buyers in these informal supply chains: *asymmetry in demand* and *costly underage and overage*.

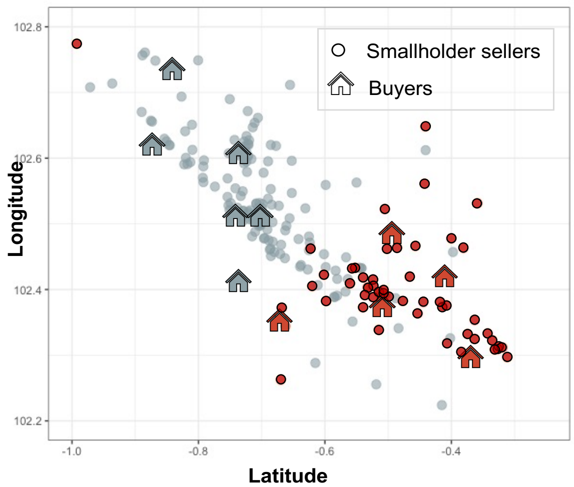

(a) Sellers trade more with nearby buyers

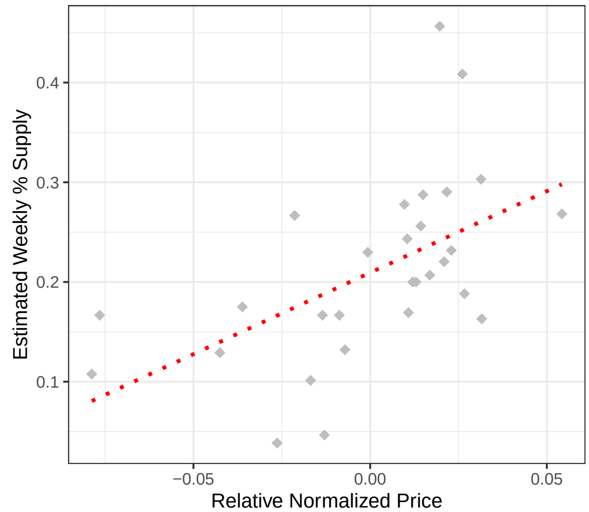

(b) Supply is price elastic and uncertain

Figure 1 Figure (a) depicts the location of buyers and smallholder sellers, colored based on frequency of trade: the grey sellers trade more with buyers in grey and the red sellers trade more with buyers in red. Figure (b) shows Buyers can attract more supply (y-axis) by setting higher bidding prices (x-axis). Appendix A.1 provides more details on how Figure (b) is generated. (Data Source: PemPem Inc.)

Asymmetry in demand arises from buyers' contractual obligations, which are common across supply chains including commodity supply chains supplied by smallholder farmers. According to field interviews with buyers of smallholder palm fruit, a typical buyer (a palm oil mill) has contractual commitments to deliver to a downstream buyer at regular time intervals, with purchase orders specifying quantity, price and delivery dates. Such contractual obligations appear to be an important factor in buyers' pricing decisions, as interviewees stated that *"In practice, some mills peak the price when they need to get more [supply] to meet their contractual obligations."* We refer to an *asymmetry* in demand because these contractual obligations typically vary across buyers, for a multitude of reasons such as processing capacity, asynchronized timeline and relationships with their downstream buyers. Using buyers' prices posted on *PemPem*, we indeed identify frequent asynchronous price adjustments for smallholder palm fruit that are not explained by the prices of the underlying commodity (crude palm oil) (see Figure 2). Because the asymmetry in demand is driven by contractual obligations, we henceforth also refer to it as *contractual imbalance*.

Underage and overage costs also arise under contractual obligations. The strategy of peaking prices to attract supply implies the existence of underage costs that incentivize meeting the target quantity at the cost of narrower profit margins. Meanwhile, overage costs exist due to factors such as holding costs (including financing costs), waste, or spoilage. Costs associated with underage and overage are further exacerbated by prevailing supply uncertainty in such environments, driven by

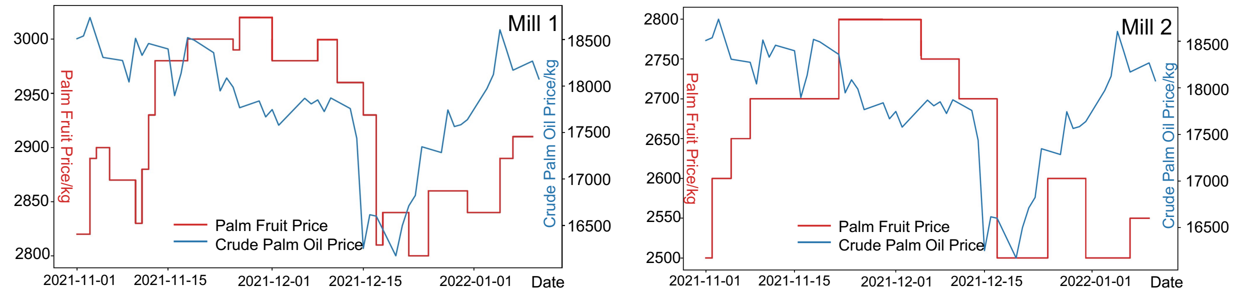

Figure 2 **Asynchronous palm fruit price adjustments by two unique palm fruit buyers vs. crude palm oil price (Source: PemPem Inc.)**

both yield uncertainty and sellers' noisy selling decisions due to informal business relationships (see, *e.g.*,Borodin et al. 2016, Gokarn and Kuthambalayan 2017, Knapp and van der Heijden 2018).

We fully characterize the pure strategy competitive equilibrium and illustrate that increased price transparency impacts market welfare via two key mechanisms. The first is *price competition*. Increasing transparency enables more sellers to seek out a better price, thus increasing supply elasticity and introducing or escalating price competition among buyers. Sellers benefit from the improved prices. The second is *demand signaling*. Price transparency enables buyers with contractual imbalances to leverage differentiated pricing to signal unequal demand, thereby reducing demand-supply mismatch and associated underage or overage costs. This benefits buyers. For both buyers and sellers to be better off, there needs to be both price competition benefiting sellers, and significant benefits from demand signaling for buyers that outweigh losses from price competition.

We identify conditions that guarantee two-sided benefits driven by price competition and demand signaling. Our results reveal that two-sided benefits are most likely to arise in markets with two defining characteristics: high overage costs that make excessive supply undesirable, and high contractual imbalance that dominates random supply variations in the market. In addition, the effect of demand signaling is enhanced in low supply uncertainty environments, where lower levels of contractual imbalance can also lead to two-sided benefits. Importantly, we also show that achieving two-sided benefits in any market further requires implementing a well-chosen level of price transparency by the platform designer: the level of transparency must be high enough to create price competition, and not too much such that heightened competition hurts buyers' profit.

To address the concern that high transparency makes cooperation easier and incentivizes price collusion (see, e.g. Stigler 1964, Green and Porter 1984), we further examine how our results are affected by possible collusion. We find that although colluding buyers might price below the competitive market price, the existence of demand signaling can still lead to two-sided benefits.

Our findings offer important managerial insights for information platforms to effectively identify target markets and ensure sustainable information provision. Our key recommendations are the following (see also Table 1 of Section 7.1 for a summary).

Target markets. Platforms seeking to achieve two-sided benefits should focus on markets where buyers face high overage costs and significant demand asymmetry. Additionally, low supply uncertainty enhances platform performance; if supply uncertainty is high, platforms can mitigate this by fostering a low-uncertainty environment for the buyers. A platform can do so, for example, by integrating centralized matching into the platform–an approach that *PemPem* is now pursuing.

Information provision strategy. To determine the optimal level of transparency, platforms should assess the intensity of price competition in the market. When high underage costs and low market differentiation drive intense price competition, partial transparency is recommended to protect buyer profits. This can be achieved by limiting disclosed information uniformly across sellers overtime (e.g. by offering a subset of data or aggregated data) to maintain fairness. In low-competition markets (low underage costs, high market differentiation), full transparency is ideal as it maximizes market efficiency. The corresponding information disclosure polices may be region-sensitive and may require active re-balancing overtime.

The remainder of this paper is structured as follows. Section 1.1 describes related literature and Section 2 describes the model. The pure strategy competitive equilibrium is characterized in Section 3, and Section 4 presents our main results under the competitive equilibrium. Section 5 presents our results under collusive equilibria. Section 6 extends the model and our main result to more general settings. In Section 7, we discuss general managerial insights and specific recommendations using model calibrated for *PemPem*'s operations. Finally, Section 8 provides our concluding remarks. All proofs can be found in Appendix B-E.

1.1. Related Literature

An extensive empirical literature has examined the impact of transparency on price and welfare in the developing world. In general, there is no empirical consensus on whether or when there are welfare benefits to both sides of the market. Goyal (2010) observes that improving price transparency leads to much more competitive market prices, resulting in higher supplier welfare but a loss in buyer profits. Similarly, Svensson and Yanagizawa (2009) and Abraham (2006) focus on one-sided benefits for the price-searching supplier. Meanwhile, Fafchamps and Minten (2012), Futch and McIntosh (2009), Aker and Fafchamps (2015) find no such welfare impacts. Levi et al. (2020) identifies improved welfare for suppliers of some products but no significant impact for the rest. Our work offers possible explanations for the variations in empirical findings.

The market conditions in our model are most closely related to those in the seminal empirical paper by Jensen (2007), which examines a setting with supply uncertainty, demand asymmetry, geographical differentiation, and high overage costs due to perishable goods. Using field data, the author observes two-sided welfare benefits from price transparency due to better demand-supply matching and reduced waste. Although a simple model is proposed to explain the empirical findings, it only explains the suppliers' behavior and one-sided benefits to the suppliers. The author remarks, "how the net welfare gain is shared between the two groups, and whether, in fact, one group gains while the other loses... is *a priori* ambiguous." In contrast, our proposed model incorporates strategic behavior from both sides and identifies market conditions and levels of price transparency needed to achieve two-sided benefits.

In addition, many of those empirical studies are also interested in how new information technologies impact the "law of one price" in supply chains of the developing world. While the general finding is that improved transparency reduces price dispersion in the market (see, e.g. Aker and Fafchamps 2015, Aker 2010, Jensen 2007, Parker et al. 2016, Aker 2008), many also reported that price dispersion may persist after the introduction of information technologies (see, e.g. Parker et al. 2016, Aker 2008). Explanations put forward to explain such persistent price dispersion include the price searcher's failure to compare prices (Baye and Morgan 2001), obfuscation strategies (Ellison and Ellison 2009) and more. We contribute to this line of literature by providing a theoretical explanation for how price dispersion may persist due to the presence of contractual imbalance.

Existing theoretical literature on price transparency provides mixed results on the price and welfare impact of improved transparency. Many contributions in the economics literature argue that improvements in price transparency on the price-searching side lead to more competitive pricing. These arguments usually build on the results of the consumer search literature of the 1970s and 1980s, such as Salop and Stiglitz (1977), Varian (1980), Burdett and Judd (1983), Stahl (1989) and many others. Price competition in turn benefits the price searchers. However, the impact of price transparency can be ambiguous when price collusion is considered. In differentiated Hotelling markets, Schultz (2005), Rasch and Herre (2013) and Schultz (2017) showed that whether improved transparency on the side of price searchers makes tacit collusion more or less difficult depends on the level of differentiation in the markets. Meanwhile, improved price transparency on the side of the price-setters alone is mostly viewed as anti-competitive as it facilitates tacit collusion (see, e.g. Stigler 1964, Green and Porter 1984).

None of the above theoretical research mentions the possibility of two-sided benefits, and none considers the price-setters' broader operations under uncertainty or contractual obligations with shortage and holding costs. To incorporate the buyers' operations in our model, we draw on techniques from a separate line of revenue management literature on price-setting newsvendor models (see Petruzzi and Dada (1999) for a comprehensive review). These newsvendor models focus on the behavior of price-setters under uncertainty and known quantity targets. Similar to our setting, many study newsvendor *buyers* facing random supply (see, e.g. Keren 2009, Chen and Liu 2021), and some also consider duopolistic competition among the newsvendor firms and provide analysis of the Nash equilibria (see, e.g. Parlar 1988, Lippman and McCardle 1997, Tsay and Agrawal 2000, Güler et al. 2018). However, the existing newsvendor models do not explicitly model price search of the other side of the market under limited information. In contrast, our work contributes to the literature by explicitly modeling both the newsvendor behavior of the competing (or colluding) buyers and the price searching behavior of the suppliers.

More recently, a growing OM literature studies information sharing for poverty alleviation in global supply chains. This line of research predominantly focuses on improving welfare of economically vulnerable suppliers through information provision. Sodhi and Tang (2014) and Chen and Tang (2015) investigate the benefits of information in reducing the search cost, uncertainty and making better selling decisions by suppliers in developing countries. Chen et al. (2015), Xiao et al. (2020) study the issue of incentivizing information sharing among smallholder farmers. Zhou et al. (2021), Liao et al. (2019) look at design of optimal information provision to maximize farmers' welfare by a social planner. He et al. (2018) studies formation of information-sharing coalitions by farmers under information provision. We contribute to this line of research by considering information provision that achieves two-sided benefits in such contexts.

Lastly, our work contributes more broadly to the growing literature on online platforms and information design. Many have studied information design questions using a Bayesian persuasion framework (see Bergemann and Morris (2019) for a comprehensive review). Others also examined optimal information elicitation (e.g. Palley and Soll (2019)) and optimal information provision for optimal learning (e.g. Kremer et al. (2014), Papanastasiou et al. (2018)). To our knowledge, our work is the first to consider two-sided benefits as a design objective of information design.

2. Model

We consider a Hotelling market with a continuum of smallholder sellers located uniformly across $[0,1]$. The total supply from the sellers every period is random, with constant expectation normalized to 2. Two buyers, located at $x = 0$ and $x = 1$, play a multi-period duopoly game. In each period, each buyer sets a per-unit price p at which he buys produce from the smallholder sellers. Both buyers aim to meet their own pre-committed volume contract with a downstream buyer.

2.1. Seller Utilities

A seller who sells to buyer $i \in \{0,1\}$ located distance $d \in [0,1]$ away at a price p_i receives utility $p_i - td$ per unit of produce. The parameter $t > 0$ is the transportation cost per unit distance and reflects the level of geographical differentiation of the market. Since t affects decision-making only for price searchers, it can be more generally interpreted as the level of differentiation between buyers that leads to seller stickiness or loyalty with a main buyer.

A seller located at $x \in [0,1]$ has a reservation price $\underline{p}(x) \geqslant 0$. That is, she desires to sell all of her produce if the price is at or above $\underline{p}(x)$, and will sell none of it below $\underline{p}(x)$. Generally, different sellers may have different valuations for $\underline{p}(x)$, and the same seller can have different reservation prices at each buyer due to unequal transportation costs. For analytic tractability, we will focus on a simplified setting where all sellers share an identical reservation price $\underline{p}$ at both buyers.

Assumption 1. *All sellers have uniform reservation price $\underline{p}$, such that they always sell all of their produce at or above $\underline{p}$, and sell none below $\underline{p}$.*

Since all sellers have reservation price at $\underline{p}$, both buyers must price at or above $\underline{p}$ to get non-zero supply.

It is worth noting that while this assumption improves tractability of our analysis, it does not drive our key findings. In Appendix A.4, we consider a more realistic setting where each seller sets lower reservation price at the buyer closer to him compared to the buyer further away due to different transportation costs. We show that this adjustment has limited impact on our qualitative insights under either competitive or collusive pricing.

2.2. Supply Function

Each buyer has informal relationships with a set of loyal sellers. Without access to an information platform, sellers are uninformed about the competing prices and will always sell to the buyer they are loyal to. Since the cost of transportation is high, sellers can only visit one buyer per period.[1] In expectation, half of the total supply is from sellers loyal to buyer 0 and the other half is from sellers loyal to buyer 1, so that market power is balanced between the two buyers. We do not require any assumptions on the distribution of loyalty on the Hotelling street; one intuitive distribution is that each seller sells to the buyer that is closer to her.

With a platform providing price transparency, a fraction $\lambda \in (0,1]$ of the sellers becomes *informed* about both prices p_0, p_1 and their distances to both buyers, while the rest remain uninformed. informed sellers are distributed uniformly across $[0,1]$. An informed seller always sells to the utility-maximizing buyer, and is indifferent between selling to either buyer if located at $x = \frac{1}{2} + \frac{p_0 - p_1}{2t}$. The parameter λ is our measure of price transparency in the market, and it can be viewed as a design lever that the platform tunes through its information disclosure strategy towards sellers. We discuss how this may be implemented in Section 7.3.

[1] For instance, fishermen in South India studied in Jensen (2007) can only visit one market per day due to high transportation costs and the limited duration of the market. The same is also observed in the smallholder palm fruit market that *Pempem* operates in.

The expected supply of buyer i, given his price $p_i \geqslant \underline{p}$ and his competitor's price $p_j \geqslant \underline{p}$, is denoted by $y(p_i, p_j, \lambda)$. It is the sum of i) supply from the $(1-\lambda)$ fraction of uninformed loyal sellers and ii) that from the informed sellers who receive higher utilities by selling to i:

$$y(p_i, p_j, \lambda) = (1-\lambda) + 2\lambda \cdot \min\left\{ \left(\frac{1}{2} + \frac{p_i - p_j}{2t} \right)^+, 1 \right\}. \tag{1}$$

The expected supply of buyer j, $y(p_j, p_i, \lambda)$, is symmetrically defined, and it is straightforward to verify that $y(p_i, p_j, \lambda) + y(p_j, p_i, \lambda) = 2$.

Each buyer faces uncertain supply due to factors such as varying yield, harvest schedule, and selling decisions from sellers, which are not observed by buyers until the end of each period. We thus model buyer i's supply as $y(p_i, p_j, \lambda) + \epsilon_i$, where ϵ_i is a continuous random variable with zero mean and cumulative distribution function

$$\epsilon_i \sim F(\cdot; \sigma).$$

Here, F is parameterized by the scale parameter $\sigma > 0$, which measures the level of supply uncertainty in a given market. Greater values of σ correspond to higher variances and greater supply uncertainty. Formally, $F(x;\sigma) = F(x/\sigma; 1)$, and $f(x;\sigma) = \frac{1}{\sigma} f(x/\sigma; 1)$. Note that we assume $\sigma > 0$ such that the distribution is always well-defined.[2] Such an assumption is natural in the setting of smallholder supply chains, where farmers' yield and decisions are intrinsically uncertain.

Through the uncertainty term ϵ_i, we aim to capture the two most salient sources of uncertainty in smallholder supply chains: yield uncertainty of individual sellers and randomness in individual sellers' selling decisions. Figure 13 in Appendix A.3 demonstrate the presence of both factors in our partner platform's market. In Appendix A.3, we provide models that explicitly capture these two sources of uncertainty, and show that they lead to various forms of supply distributions faced by buyers.

For most of our analysis (Section 3-5), we will focus on additive supply uncertainty distributions that are independent of the expected supply, y, and of transparency λ. This simple, separable form of uncertainty is common in the newsvendor literature (see, e.g., Petruzzi and Dada (1999)). It allows us to obtain tractable characterization of the equilibrium, isolate the effect of uncertainty on market dynamics and fully understand its welfare implications. We then establish in Section 6.1 that our key insights extend to more realistic uncertainty models, described in Appendix A.3, which better capture real-world supply uncertainties in smallholder markets.

The two buyers face the same marginal distribution of supply uncertainty, i.e. $\epsilon_i, \epsilon_j \sim F$, and we write ϵ as shorthand for either ϵ_i or ϵ_j. We make no assumption on the dependence between

[2] A model with no supply uncertainty allows for an infinite number of pure strategy equilibria and requires separate analysis.

ϵ_i, ϵ_j. For example, when uncertainty arises due to yield uncertainty, both ϵ_i, ϵ_j may be independent or positively correlated (see Model 1(a) and 1(b) in Appendix A.3); when uncertainty arises due to sellers switching buyers, realized uncertainty of both buyers can be negatively correlated (see Model 2 in Appendix A.3). While ϵ_i, ϵ_j can be either independent or dependent, this does not affect buyers' pricing decisions because each buyer's focus is solely on their own realized supply, which determines their profit.

Finally, for tractability, we assume that the density function of ϵ is symmetric around 0 and is strictly positive in the vicinity of $x = 0$.

Assumption 2. *The density function is symmetric around 0, i.e.,* $f(x;\sigma) = f(-x;\sigma)\ \forall \sigma, x \in \mathbb{R}$. *In addition,* $F(x;\sigma) > \frac{1}{2}$ *for all* $x > 0$ *and* $F(x;\sigma) < \frac{1}{2}$ *for all* $x < 0$.

2.3. Buyer Demand

Under contractual obligations, both buyers aim to fulfill regular, recurring procurement targets with unequal quantities. Each period, depending on timing of contracts and required volumes, the buyers may either be in a high-demand state with target volume Q_H, or in a low-demand state with target volume Q_L. We define $\Delta = \frac{1}{2}(Q_H - Q_L) \geqslant 0$ as the *contractual imbalance*, which is our measure of the level of demand asymmetry in the market.

At the end of each period, buyers will deliver all procured supply to receive a per-unit normalized revenue of 1. Given the target volume $Q_i \in \{Q_H, Q_L\}$, a buyer also incurs a fixed *underage cost*, $\gamma \geqslant 0$, for every unit of supply shortage below Q_i, and a fixed *overage cost*, $h \geqslant 0$, for every unit of supply surplus above Q_i. We assume that at least one of γ, h is strictly positive, i.e., $\gamma + h > 0$.

In practice, there may be multiple periods between each required delivery, and contractual imbalance can arise from asynchronized delivery timelines. For example, in a given period, some buyers may need to source a large quantity to fulfill a next-day delivery target (high demand state), while others are under less pressure (low demand state). Since buyers are often myopic in their pricing decisions, we simplify the setting by standardizing delivery timelines to focus on single-period delivery with varying quantities. In addition, since delivery quantities and timelines are typically negotiated in long-term contracts on a seasonal or yearly basis, buyers have limited control over their per-period demand once their contractual obligations are set.

For most of our analysis, we will consider scenarios where buyers are in opposite demand states every period, and $Q_H + Q_L = 2$ so that total supply and total demand are balanced in expectation. Over time, each buyer has an equal probability of having high or low demand.

We focus on opposite demand states because, when buyers have identical demand, there is no contractual imbalance and thus no advantage from demand signaling under improved transparency. Therefore, *strict* two-sided benefits only occur when buyers have unequal demand, and, in our simplified setting, in opposite demand states. In Appendix A.5, we discuss the general case where demand states are arbitrarily correlated and show that two-sided benefits persist on average over time periods as long as buyers are infrequently simultaneously in the high demand state. However, buyers only *strictly* benefit during periods of unequal demand.

The simplification $Q_H + Q_L = 2$ allows us to obtain a tractable characterization of the equilibrium and fully capture the underlying mechanisms of welfare impact, focusing on markets where supply-demand mismatch is a main source of inefficiency instead of overall supply or demand imbalance. Since Q_H, Q_L are determined by buyers' contractual obligations, it is logical to assume that total target quantities align with overall market supply. We note that our main result also generalizes to arbitrary levels of Q_i, Q_j (see Appendix A.5). In this more general case, however, the exact form of equilibrium prices becomes intractable.

2.4. Sequence of Events

Each period, the following sequence of events takes place.

1. Buyers i, j observe their realized demand, Q_i, Q_j, to meet their pre-committed volume contract with a downstream buyer.
2. The buyers simultaneously set their per-unit prices p_i and p_j with the target of procuring Q_i and Q_j units of produce respectively.
3. Informed sellers observe both prices, (p_i, p_j), while uninformed sellers observe one of the two prices. Sellers harvest supply and make selling decisions (Section 2.2).
4. Each buyer purchases his realized supply $y(p_i, p_j, \lambda) + \epsilon_i$, paying the unit price p_i.
5. Buyers deliver all realized supply and receive a per-unit revenue of 1, while incurring a per-unit underage cost γ for any shortage or a per-unit overage cost h for any surplus.

It is worth noting that buyers in this setting always accept all realized supply, even when the overage cost is larger than the profit margin, i.e., $h > 1 - p_i$. In practice, buyers may do this to maintain long-term relationships with sellers, or when it is difficult to coordinate procurement in decentralized locations[3]. Alternatively, however, buyers may choose to turn down unwanted supply, leading to waste in the market, as observed in the fisherman's market in Jensen (2007). Since this alternative scenario is less tractable and yields qualitatively similar takeaways, we analyze it as an extension in Section 6.2 and show that our main results still hold.

2.5. Buyer Profits

Following the analysis by Petruzzi and Dada (1999), we let $s_i(p_i, p_j, \lambda) = y(p_i, p_j, \lambda) - Q_i$ be the safety stock acquired by buyer i. The realized profit of buyer i in each period is given by

$$[y(p_i, p_j, \lambda) + \epsilon](1 - p_i) - \gamma[s_i(p_i, p_j, \lambda) + \epsilon]^- - h[s_i(p_i, p_j, \lambda) + \epsilon]^+, \tag{2}$$

[3] For instance, in the palm fruit market, each mill has multiple decentralized drop off locations, and it is difficult to coordinate among all agents during each period.

where $x^+ = \max\{x, 0\}, x^- = -\min\{x, 0\}$ denote the positive and negative part of x respectively. The single-period expected profit $\pi_i(p_i, p_j, \lambda)$ for buyer i at the time of pricing can be written as

$$\pi_i(p_i, p_j, \lambda) = \Psi(p_i, p_j, \lambda) - L_i(p_i, p_j, \lambda), \text{ where} \tag{3}$$

$$\Psi(p_i, p_j, \lambda) = y(p_i, p_j, \lambda)(1 - p_i),$$

$$L_i(p_i, p_j, \lambda) = \gamma \mathbb{E}\left[(s_i(p_i, p_j, \lambda) + \epsilon)^-\right] + h\mathbb{E}\left[(s_i(p_i, p_j, \lambda) + \epsilon)^+\right].$$

Following the terminology of newsvendor models, we refer to $\Psi(p_i, p_j, \lambda)$ as the *riskless profit* and $L_i(p_i, p_j, \lambda)$ as the *expected loss.*

Additionally, we note that this model assumes buyers cannot engage in price discrimination, i.e., charging the reservation price for uninformed sellers while setting a different price on the platform. We provide a brief discussion of how price discrimination can affect our results in Appendix A.6.

3. Competitive Equilibria in the Single Stage Game

We start by considering the single stage game where each buyer is myopic in pricing and maximizes his expected profit for the current period. As buyers 0 and 1 are symmetric in their pricing strategies, we will use p_H, p_L to denote the prices set by the buyer in the high and low demand state respectively, and use the subscripts $\{i, j\} = \{H, L\}$. For conciseness, we will also denote the market prices using $\mathbf{p} = (p_H, p_L)$ if there is no ambiguity in the given context.

In each period, each buyer observes his demand state i and decides on his best price response, p_i^{BR}, to the competitor's price p_j by solving the following problem:

$$p_i^{BR}(p_j, \lambda) = \arg\max_{p_i \geqslant \underline{p}} \pi_i(p_i, p_j, \lambda), \quad \text{for} \quad i \in \{H, L\}.$$

In the initial market where $\lambda = 0$, both buyers price at the monopoly price, i.e., sellers' reservation price, $\mathbf{p} = \underline{\mathbf{p}} \equiv (\underline{p}, \underline{p})$. When $\lambda > 0$, the pure Bayes-Nash equilibrium for the single stage game is characterized as follows.

Theorem 1. *There exists $\hat{t}$ such that, for any $\lambda \in (0, 1]$ and $t \geqslant \hat{t}$, there is a unique pure-strategy Bayes-Nash equilibrium, and the equilibrium prices $\mathbf{p}^*(\lambda)(\lambda) = (p_H^*(\lambda), p_L^*(\lambda))$ satisfy*

$$p_H^*(\lambda) = p_L^*(\lambda) = \underline{p} \qquad \textit{if} \quad \lambda \leqslant \frac{t}{(\gamma + h)F(\Delta; \sigma) - h + (1 - \underline{p})}, \tag{4}$$

$$p_H^*(\lambda) > \underline{p}, p_L^*(\lambda) \geqslant \underline{p} \quad \textit{otherwise.} \tag{5}$$

Moreover, if $p_H^*(\lambda) > \underline{p}, p_L^*(\lambda) > \underline{p}$, the average price $\bar{p}^*(\lambda)$ and price dispersion $p_\Delta^*(\lambda)$ satisfy

$$\bar{p}^*(\lambda) := \frac{1}{2}\left[p_H^*(\lambda) + p_L^*(\lambda)\right] = 1 + \frac{\gamma - h}{2} - \frac{t}{\lambda}, \tag{6}$$

$$p_\Delta^*(\lambda) := p_H^*(\lambda) - p_L^*(\lambda) = \frac{\gamma + h}{3}\left[1 - 2F\left(\frac{\lambda}{t}p_\Delta^*(\lambda) - \Delta; \sigma\right)\right]. \tag{7}$$

Figure 3 illustrates the equilibrium prices and price dispersion as a function of λ described by the theorem with uniformly distributed supply uncertainty, $\epsilon \sim U[-\sigma,\sigma]$. A full characterization of equilibrium prices can be found in Appendix C.

Theorem 1 demonstrates the dual effects of price competition and demand signaling.

Price competition. When transparency is low, both buyers keep their prices at the reservation price (Equation (4)). They start competing on price only under sufficiently high transparency, $\lambda > \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$, where at least one of them prices above $\underline{p}$. Above $\underline{p}$, the average market price $\bar{p}^*(\lambda)$ is strictly increasing in λ (Equation (6)). Price competition escalates with increasing transparency because more informed sellers seek out a better price, making supply more elastic. Such an effect is widely observed in standard price search models from existing literature (see, for example, Salop and Stiglitz (1977), Varian (1980), Burdett and Judd (1983), Stahl (1989)). When $\gamma = h = 0$, the equilibrium prices in our model simplify to $p_H^*(\lambda) = p_L^*(\lambda) = \max\{\underline{p}, 1-\frac{t}{\lambda}\}$, recovering the standard characterization obtained in Schultz (2005) where sellers are price setters. Unique to our context, however, is the impact of underage/overage costs on the average price (Equation (6)). Higher underage costs increase buyer competition for additional safety stock to avoid potential shortage; higher overage costs, on the other hand, discourage competition to avoid unwanted surplus. In markets with very high overage costs relative to underage costs, price competition may thus not occur even at near full transparency.

Demand signaling. As competitive pricing drives prices above the reservation price, price dispersion arises when $\gamma + h > 0$ and $\Delta > 0$ (Equation (7)). This dispersion is driven by costly overage/underage and contractual imbalances. Under sufficient transparency, the buyers can signal the differences in their contractual demand through pricing, $p_H^* > p_L^*$, and attract supply accordingly to reduce overage and underage costs; with insufficient transparency, differential pricing would not make much of a difference and so price dispersion is zero. This leads to an initial increase in $p_\Delta^*(\lambda)$ from zero when sufficient transparency is introduced. After the initial increase, transparency can have a non-monotonic effect on price dispersion. Moreover, price dispersion is not only affected by transparency λ but also by supply uncertainty σ in the market. Proposition 1 describes the relationships between transparency, supply uncertainty and price dispersion.

Proposition 1. *For any $\Delta > 0$ and $\gamma + h > 0$, price dispersion $p_\Delta^*(\lambda)$ has the following properties:*

1. *When both buyers' prices are above the reservation price, greater transparency always reduces price dispersion in the market, i.e., $\frac{dp_\Delta^*(\lambda)}{d\lambda} \leqslant 0$ for any λ such that $p_H^*(\lambda), p_L^*(\lambda) > \underline{p}$.*
2. *Greater supply uncertainty always reduces price dispersion in the market, i.e., $\frac{\partial p_\Delta^*(\lambda)}{\partial\sigma} \leqslant 0$ for any fixed λ such that $p_\Delta^*(\lambda) > 0$.*

Moreover, both inequalities are strict if the density function is strictly positive at $\epsilon = \Delta - \frac{\lambda}{t}p_\Delta^*(\lambda)$, i.e., $f(\Delta - \frac{\lambda}{t}p_\Delta^*(\lambda);\sigma) > 0$.

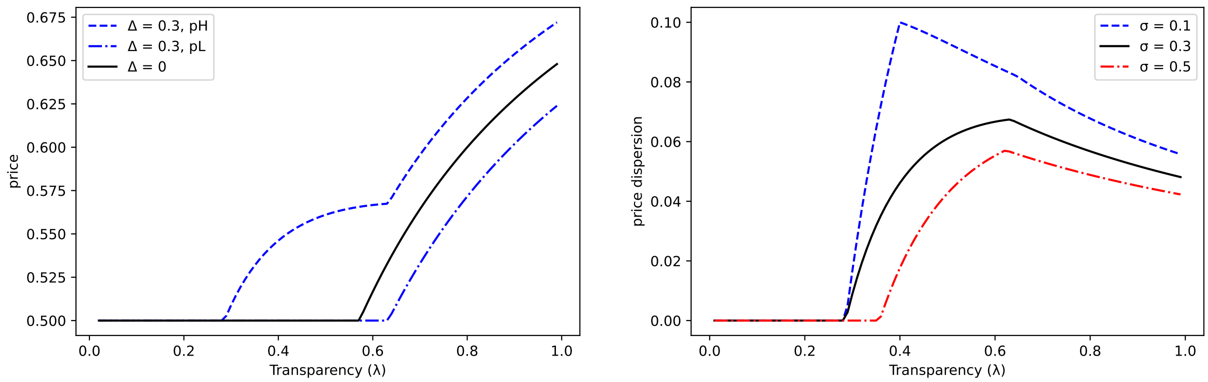

(a) $p_H^*(\lambda)$, $p_L^*(\lambda)$ v.s. λ for $\sigma = 0.3$ (b) $p_\Delta^*(\lambda)$ v.s. λ for $\Delta = 0.3$

Figure 3 Equilibrium prices and price dispersion in a market with uniformly distributed supply uncertainty, $\epsilon \sim U[-\sigma, \sigma]$**. The other parameters are** $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2$**.**

Proposition 1 states that once both buyers price above the reservation price, the equilibrium price dispersion $p_\Delta^*(\lambda)$ decreases with increasing transparency. Equilibrium price dispersion also decreases with greater supply uncertainty under fixed transparency. This relationship is illustrated in Figure 3b, where the equilibrium price dispersion is plotted for different values of σ. One way to interpret this relationship is that increasing supply elasticity reduces the level of price difference needed to achieve the desired adjustment in expected supply; meanwhile, increasing supply uncertainty reduces buyers' incentives to use price signals. The first part of the result agrees with empirical findings that transparency generally reduces price dispersion (see, e.g. Aker and Fafchamps 2015, Aker 2010, Jensen 2007, Parker et al. 2016, Aker 2008). Moreover, the second part further reveals that higher supply uncertainty weakens the effect of demand signaling in the market. This is because the buyers differentiate their prices so that the buyer with high (low) demand procures more (less) supply in expectation; however, in doing so, the high (low) demand buyer risks over-procuring (under-procuring). When supply is highly uncertain, such tail risks are significant, reducing the benefits of raising (or lowering) prices in the high (or low) demand state.

Remarks. A minimal level of geographical differentiation ($\hat{t}$) is required in Theorem 1 for two purposes. First, the existence of a pure Nash equilibrium is not guaranteed if geographical differentiation t is too low. This is because, as in Schultz (2005) and Varian (1980), the competitive prices given in (6) and (7) can be so high that it may be better for a buyer to lower his price to $\underline{p}$ and only buy from the $1-\lambda$ uninformed sellers, i.e., $\pi_i(p_i^*, p_j^*, \lambda) < \pi_i(\underline{p}, p_j^*, \lambda)$. When this happens, only mixed strategy equilibria exist. While Schultz (2005) and Varian (1980) characterize mixed strategy equilibria in symmetric markets with $\Delta = \gamma = h = 0$, in general, such mixed strategy equilibria do not admit a closed-form characterization (Schultz 2005). Second, if t is too low under asymmetric demand ($\Delta > 0$), price dispersion characterized in Equation (7) may exceed t. When this happens, price dispersion is capped at t, leading to an equilibrium where the use of price signals is constrained by limited geographical differentiation.[4] Note that the qualitative insights in this case are similar. We characterize $\hat{t}$ in Appendix C.4, and focus on markets with $t \geqslant \hat{t}$.

4. Welfare Under Competitive Equilibrium

We now present our main results on the welfare impact of transparency under competitive equilibrium. Our results reveal how price competition and demand signaling act together under contractual imbalance and costly underage or overage to drive welfare changes, and characterize sufficient conditions for strong Pareto improvement.

4.1. Measuring Welfare and Market Efficiency

We define buyer and seller welfare using the expected profit for each buyer ($\bar{\pi}$) and the expected utility for each seller ($\bar{u}$) respectively. Recall that over multiple time periods, each buyer has an equal chance of being in high or low demand every period; the bar notation indicates that a quantity is being taken in expectation over buyer demand states as well as over uncertainty in supply.

Given any set of market prices $\mathbf{p} = (p_H, p_L)$, each buyer's single period expected profit in high and low demand can be denoted using the short forms $\pi_H(\mathbf{p}, \lambda)$ and $\pi_L(\mathbf{p}, \lambda)$ respectively. Expected buyer welfare, denoted by $\bar{\pi}(\mathbf{p}, \lambda)$, is thus the expected profit averaged over demand states,

$$\bar{\pi}(\mathbf{p}, \lambda) = \frac{1}{2}[\pi_H(\mathbf{p}, \lambda) + \pi_L(\mathbf{p}, \lambda)]. \tag{8}$$

For an uninformed seller, her per-unit expected utility is $\bar{u}^U(\mathbf{p}, \lambda, d) = \frac{1}{2}[p_H + p_L] - td$, where d is her distance to the buyer that she is loyal to and superscript U denotes the uninformed state. After a seller becomes informed, her per-unit expected utility becomes

$$\bar{u}^I(\mathbf{p}, \lambda, d) = \frac{1}{2}\left(\max[p_H - td, p_L - t(1-d)] + \max[p_L - td, p_H - t(1-d)]\right). \tag{9}$$

In the initial market where $\lambda = 0$, the buyers always price at sellers' reservation price, $\underline{\mathbf{p}} \equiv (\underline{p}, \underline{p})$. We say that a level of transparency λ is *strongly Pareto improving* under $\mathbf{p}$ if the welfare of all sellers and both buyers strictly increases compared to in the initial market, i.e.,

$$\bar{\pi}(\mathbf{p}, \lambda) > \bar{\pi}(\underline{\mathbf{p}}, 0), \quad \text{and} \quad \bar{u}^I(\mathbf{p}, \lambda, d) \geqslant \bar{u}^U(\mathbf{p}, \lambda, d) > \bar{u}^U(\underline{\mathbf{p}}, 0, d) \quad \forall d \in [0, 1]. \tag{10}$$

To better understand drivers of market welfare, we introduce the notion of *market efficiency*, measured by the expected level of supply-demand mismatch in the market. Recall from Section 2.3 that for given market prices $\mathbf{p} = (p_H, p_L)$, we use $s_H(p_H, p_L, \lambda)$ and $s_L(p_L, p_H, \lambda)$ to denote the

[4] The equilibrium price dispersion is capped at t because when buyer i prices above $p_j + t$, he attracts no additional expected supply (see Equation (1)). His expected profit thus always improves if p_i is lowered to match $p_j + t$.

safety stock acquired by high and low demand buyers respectively, defined as the difference between a buyer's expected supply and realized demand. In addition, $s_H(p_H, p_L, \lambda) = -s_L(p_H, p_L, \lambda)$ since total supply is equal to total demand in expectation. Consequently, we let

$$s(\mathbf{p}, \lambda) := |s_H(p_H, p_L, \lambda)| = |s_L(p_L, p_H, \lambda)| = \left| 2\lambda \min\left\{ \left(\frac{p_H - p_L}{2t} + \frac{1}{2} \right)^+, 1 \right\} - (\lambda + \Delta) \right| \tag{11}$$

be the absolute level of the mismatch between the observed demand and expected supply under prices $\mathbf{p}$ and transparency λ. We use $s(\mathbf{p}, \lambda)$ as a measure of *market efficiency*, and say that one market is more *efficient* than another if it induces a lower level of supply-demand mismatch.

4.2. When and How Strong Pareto Improvement Arises

We now discuss conditions under which strong Pareto improvement arises. We first specify the *competitive market condition*, a necessary and sufficient condition for transparency to induce higher prices and therefore improve sellers' welfare. We then analyze under what conditions increased transparency not only increases sellers' welfare but also buyers' welfare, creating two-sided benefits.

Lemma 1 (Competitive Market Condition). *All sellers are strictly better off under the competitive equilibrium* $\mathbf{p}^*(\lambda)$ *if and only if* $p_H^*(\lambda) > \underline{p}$, *or, equivalently,* $\hat{\lambda} < 1$ *and* $\lambda > \hat{\lambda}$, *where* $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma) - h + 1 - \underline{p}}$.

Lemma 1 establishes that sellers benefit from price transparency if and only if there is an increase in market price due to price competition. If transparency does not lead to competitive pricing, sellers do not benefit because buyers will keep prices at the reservation price, and there is no price variation that sellers can take advantage of.

Not all markets satisfy the first inequality, $\hat{\lambda} < 1$, or, equivalently, $(\gamma + h)F(\Delta;\sigma) - h + (1 - \underline{p}) > t$. In some markets, the overage cost h can be so high relative to the underage cost γ that the buyers have no incentive to compete for additional supply. Alternatively, the geographical differentiation t can be so high that each buyer essentially becomes a regional monopoly with no local competitor. This observation offers a potential explanation for existing empirical results such as in Fafchamps and Minten (2012), Futch and McIntosh (2009), Aker and Fafchamps (2015), Levi et al. (2020) where improved transparency had no measurable impact on prices of agricultural products. In such markets, introducing price transparency may do little to improving market prices for sellers, and achieving strong Pareto improvement will be difficult. For the rest of the section, we thus focus on markets that satisfy $\hat{\lambda} < 1$.

In such markets, the competitive market condition then requires that price transparency exceeds the said threshold level, $\hat{\lambda}$. Because sellers only benefit when price transparency exceeds this threshold, $\hat{\lambda}$ will appear in the remainder of our results.

While price competition benefits the sellers, it harms buyers' profit due to narrower profit margins. We now turn to identifying conditions under which increased transparency not only increases sellers' welfare but creates two-sided benefits for both sellers and buyers. Using Equation (3), each buyer's expected profit $\bar{\pi}(\mathbf{p},\lambda)$ can be written as

$$\bar{\pi}(\mathbf{p},\lambda) = \bar{\Psi}(\mathbf{p},\lambda) - \bar{L}(\mathbf{p},\lambda),$$

where $\bar{\Psi}$ and $\bar{L}$ represent the riskless profit and the expected loss averaged over demand states, i.e., $\bar{\Psi}(\mathbf{p},\lambda) := \frac{1}{2}[\Psi(p_H,p_L,\lambda)+\Psi(p_L,p_H,\lambda)]$, and $\bar{L}(\mathbf{p},\lambda) := \frac{1}{2}[L_H(p_H,p_L,\lambda)+L_L(p_L,p_H,\lambda)]$. Lemma 2 suggests that while the riskless profit $\bar{\Psi}$ cannot be improved compared to the initial monopolistic market, buyers may benefit from reduced expected loss, $\bar{L}$.

Lemma 2. *For any market price* $\mathbf{p}$ *and transparency level* λ*, individual buyer's expected profit has the following properties:*

1. *The riskless profit is always the highest in the initial market, i.e.,* $\bar{\Psi}(\mathbf{p},\lambda) \leqslant \bar{\Psi}(\underline{\mathbf{p}},0)$.
2. *The expected loss is reduced if and only if the expected supply-demand mismatch is reduced, i.e.,* $\bar{L}(\mathbf{p},\lambda) < \bar{L}(\underline{\mathbf{p}},0)$ *if and only if* $s(\mathbf{p},\lambda) < s(\underline{\mathbf{p}},0)$.

Lemma 2 shows that the riskless profit is maximized in the initial monopolistic market, where $\mathbf{p} = \underline{\mathbf{p}}$. Hence, in order for buyer welfare to strictly increase, there must be a reduction in the expected loss, $\bar{L}$, which can only be achieved via reducing the expected supply-demand mismatch $s(\mathbf{p},\lambda)$. Proposition 2 states that increased transparency indeed reduces supply-demand mismatch under the competitive equilibrium, thereby reducing buyers' expected loss.

Proposition 2. *For any* $\Delta > 0$*, greater transparency always strictly improves market efficiency and reduces buyers' expected loss for all* $\lambda > \hat{\lambda}$ *in the competitive equilibrium, i.e.,* $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ *and* $\frac{d\bar{L}(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ *for all* $\lambda > \hat{\lambda}$*, where* $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$.

Proposition 2 shows that higher transparency always strictly reduces supply-demand mismatch $s(\mathbf{p}^*(\lambda),\lambda)$, improving market efficiency. In turn, this reduces the buyer's expected loss $\bar{L}(\mathbf{p}^*(\lambda),\lambda)$. This is the result of demand signaling, where increased transparency enables the buyers to differentiate their prices $p_H^* > p_L^*$ under asymmetric procurement targets.

Proposition 3 further establishes that the degree of loss reduction is greater in markets with higher contractual imbalance, Δ, and lower supply uncertainty, σ.

Proposition 3. *For any fixed* $\lambda \in (\hat{\lambda},1]$*, let* $\bar{L}_\Delta(\lambda) = \bar{L}(\underline{\mathbf{p}},0) - \bar{L}(\mathbf{p}^*(\lambda),\lambda)$ *denote reduction in buyers' expected loss in the competitive equilibrium relative to* $\lambda = 0$*. Then, for any* $\Delta > 0$*,*

(i) $\bar{L}_\Delta(\lambda)$ *is increasing in the level of contractual imbalance,* $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \Delta} \geqslant 0$*, and*

(ii) $\bar{L}_\Delta(\lambda)$ *is decreasing in the level of supply uncertainty,* $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \sigma} \leqslant 0$.

Moreover, both inequalities are strict if the density function is strictly positive at $\epsilon = \Delta - \frac{\lambda}{t}p_\Delta^*(\lambda)$*, i.e.,* $f(\Delta - \frac{\lambda}{t}p_\Delta^*(\lambda);\sigma) > 0$.

While buyers facing greater demand asymmetry naturally benefit more from demand signaling, the second part of Proposition 3 highlights that greater supply uncertainty leads to less efficiency improvement and thus reduces the welfare benefits from demand signaling in equilibrium. We have seen in Proposition 1 that uncertainty discourages the buyers from using price differentiation to signal demand, lowering the equilibrium price dispersion $p^*_\Delta(\lambda)$. As a consequence, the lack of demand signaling means lower transparency-induced efficiency gains.

When benefits from demand signaling (reduced $\bar{L}$) outweigh losses from narrower profit margin due to price competition (reduced $\bar{\Psi}$) for buyers, two-sided benefits arise in the market: buyers benefit from reduced expected loss and sellers benefit from improved prices. Theorem 2 presents a set of sufficient conditions that guarantees such two-sided benefits.

Theorem 2. *In any market satisfying the competitive market condition, if $\Delta > \underline{\Delta}$ and $h \geqslant 1-\underline{p}$, then there exists a threshold level $\lambda^* \in (\hat{\lambda}, 1]$, such that any $\lambda \in (\hat{\lambda}, \lambda^*)$ is strongly Pareto improving under the competitive equilibrium. Herein, $\underline{\Delta}$ is the smallest value satisfying $F(\underline{\Delta};\sigma) = \frac{\gamma+1-\underline{p}}{\gamma+h}$, and $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$.*

Theorem 2 shows that two-sided benefits can be achieved under three joint conditions.

(1) High contractual imbalance that dominates random supply variations, $\Delta > \underline{\Delta}$. The condition on Δ demands that $F(\Delta;\sigma) \geqslant \frac{\gamma+1-\underline{p}}{\gamma+h}$. Equivalently, it requires that $\mathbb{P}[\epsilon > \Delta|\sigma]$ is sufficiently small, such that $\mathbb{P}[\epsilon > \Delta|\sigma] \leqslant \frac{h-(1-\underline{p})}{\gamma+h}$, where ϵ represents random supply variations in the market. This inequality, whenever satisfied, suggests that the degree of random supply variations, ϵ, is being dominated by the level of contractual imbalance in the market. Correspondingly, the threshold $\underline{\Delta}$ decreases as supply uncertainty σ decreases. This implies that two-sided benefits can arise in low uncertainty environments even with limited contractual imbalance. Conversely, high supply uncertainty makes two-sided benefits harder to achieve. As explained in Proposition 3, this is because both high Δ and low σ enhance the benefits of demand signaling. The following corollary formalizes that the required contractual imbalance $\underline{\Delta}$ can become arbitrarily small under sufficiently low σ.

Corollary 2.1. *In any market satisfying the competitive market condition, if $\Delta > 0$ and $h \geqslant 1-\underline{p}$, there always exists a threshold $\hat{\sigma} > 0$ such that for any $\sigma < \hat{\sigma}$, strong Pareto improvement arises under some $\lambda \in (\hat{\lambda}, 1]$, where $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$.*

(2) High overage costs, $h \geqslant 1-\underline{p}$. A high overage cost has the dual effect of reducing price competition (since additional supply becomes less desirable) and increasing the benefit of demand signaling. Both contribute to increasing buyers' profit in the competitive equilibrium. In contrast, the same is not true for high underage cost, γ. While a high γ also increases gains from demand signaling, it intensifies price competition, which leads to a non-monotonic effect on buyers' profit.

This explains why a high h is conducive for two-sided benefits, but high γ is not necessarily beneficial. It is worth noting that whenever $h \geqslant 1 - \underline{p}$, it becomes more profitable for buyers to turn down unwanted supply. While the current model assumes that the buyers always accept all realized supply, in Section 6.2 we show that our results continue to hold when buyers turn down any supply above the target quantity.

(3) Well-chosen transparency levels, $\lambda \in (\hat{\lambda}, \lambda^*)$***.*** A minimum level of transparency ($\hat{\lambda}$) is required to guarantee price competition to benefit the sellers. Transparency may also need to be capped below a maximum level (λ^*) to protect the buyers. Beyond λ^*, heightened price competition results in a transfer of welfare from buyers to sellers.

In markets satisfying the other desirable characteristics (high h, high Δ or low σ), the ideal transparency range $(\hat{\lambda}, \lambda^*)$ is mainly influenced by how competitive the buyers are, which is tied to market differentiation t and underage cost γ. This relationship is evident from the closed-form expression for the lower threshold, $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$, which increases market differentiation (small t) and decreases with costly underage (high γ). While it is more difficult to provide similar closed-form expressions for the upper bound λ^*, Proposition 4 establishes that, like $\hat{\lambda}$, λ^* typically (although not always) increases with t and decreases with γ. Additionally, Proposition 4 also establishes that stronger effects of demand signaling (large Δ, lower σ) often leads to a wider range of Pareto-improving transparency levels, $|\lambda^* - \hat{\lambda}|$.

Proposition 4. *Consider any market satisfying the conditions for the existence of two-sided benefits under Theorem 2, and let* λ^* *be the maximum transparency level such that any* $\lambda \in (\hat{\lambda}, \lambda^*)$ *is strongly Pareto improving. Then, if* $p_L^*(\lambda^*) > \underline{p}$*, the desired range of Pareto-improving transparency levels* $(\hat{\lambda}, \lambda^*)$ *has the following properties:*

(i) The width of the range $|\lambda^* - \hat{\lambda}|$ *is always increasing in* Δ *and decreasing in* σ*.*

(ii) $\hat{\lambda}, \lambda^*$ *are both always increasing in* t*, and are both decreasing in* γ *if* $\Delta \leqslant \frac{3}{4}$*.*

Thus, partial transparency tends to be preferable (see, e.g., Figure 4) when competition is intense due to low market differentiation (small t) or costly underage (high γ), whereas full transparency becomes desirable (see, e.g., Figure 5) if competition is weaker due to high market differentiation (large t) or low underage costs (small γ). Meanwhile, the range of Pareto-improving transparency levels increase with larger demand asymmetry, Δ, and lower supply uncertainty, σ. The latter effect is demonstrated in Figure 4a.

The above results suggest that to achieve two-sided benefits, the platform designer has to be strategic in both i) identifying the target market and ii) choosing the appropriate transparency level to implement. Corollary 2.1 further suggests that two-sided benefits can also be achieved in some markets by creating a low supply-uncertainty environment for buyers. We dive further into the relevant managerial implications in Section 7.

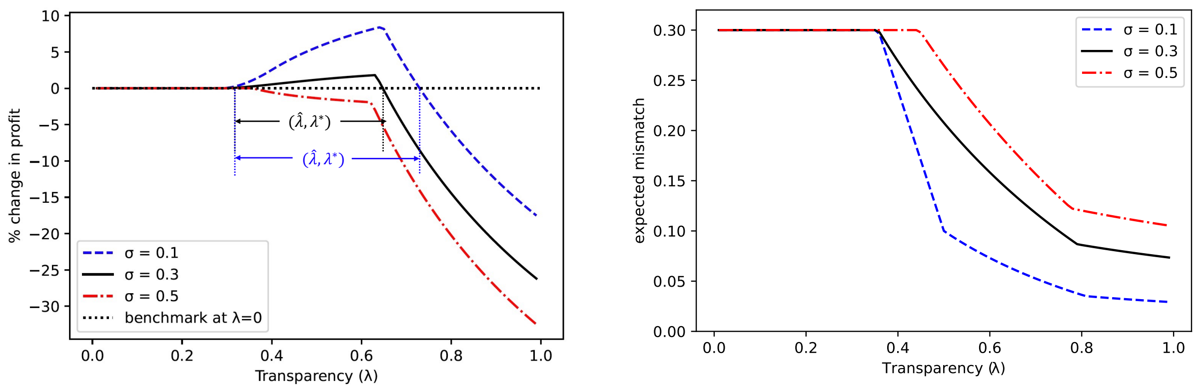

(a) $\bar{\pi}(\mathbf{p}^*(\lambda), \lambda) - \bar{\pi}(\underline{\mathbf{p}}, 0)$ v.s. λ, $t = 0.2$ (b) $s(\mathbf{p}^*(\lambda), \lambda)$ v.s. λ, $t = 0.2$

Figure 4 Two-sided benefits observed at partial transparency, under uniformly distributed supply uncertainty $\epsilon \sim [-\sigma, \sigma]$**, with** $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2$.

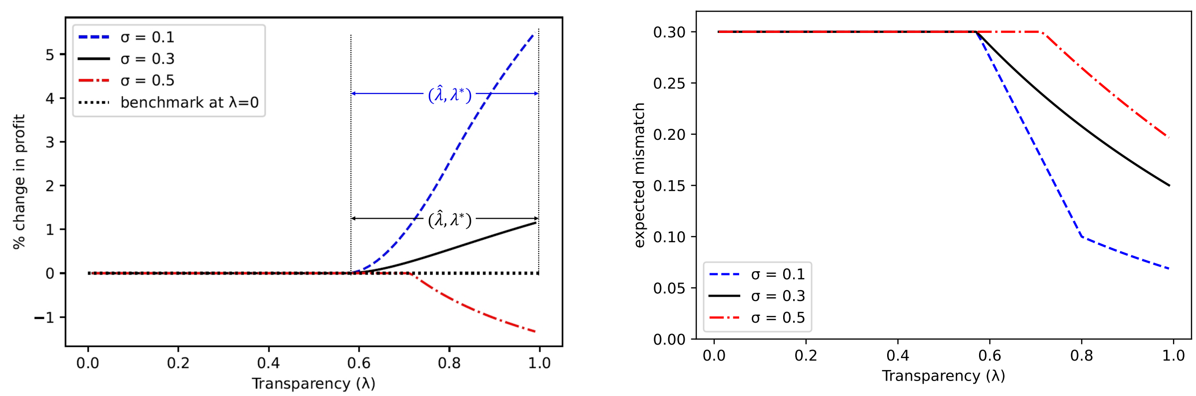

(a) $\bar{\pi}(\mathbf{p}^*(\lambda), \lambda) - \bar{\pi}(\underline{\mathbf{p}}, 0)$ v.s. λ, $t = 0.4$ (b) $s(\mathbf{p}^*(\lambda), \lambda)$ v.s. λ, $t = 0.4$

Figure 5 Two-sided benefits observed at full transparency, under uniformly distributed supply uncertainty $\epsilon \sim [-\sigma, \sigma]$**, with** $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.4$.

We end this section by pointing out a design trade-off as a consequence of our results. Proposition 2 suggests that market efficiency is maximized at full transparency, $\lambda = 1$, and so are seller utilities (Equation (6)). In contrast, Theorem 2 may recommend against too much transparency in certain markets due to excessive welfare transfer decreasing buyers' profits. The platform designer hence faces the choice between achieving two-sided benefits at an intermediate level of transparency versus maximizing efficiency and seller welfare at the highest level of transparency achievable. While this paper focuses primarily on achieving two-sided benefits, the decision will ultimately depend on the desired balance between efficiency and equity, and the specific market context.

5. Collusive Equilibria in Repeated Games

Our analysis so far established that price transparency can create two-sided benefits via the combined effects of price competition and demand signaling. However, in some markets, increasing transparency could also lead to collusion among buyers (see, e.g. Stigler 1964, Green and Porter 1984). In smallholder supply chains, price collusion is particularly concerning since it adversely impacts the welfare of the vulnerable sellers. This section examines the welfare impact of transparency under collusion and demonstrates that in the absence of price competition, increased transparency can still be beneficial to both sides due to demand signaling.

5.1. A Model for Collusion

We study a repeated game with infinitely many periods, where both buyers have the same discount factor $\delta \in [0,1]$. We consider cooperative strategies where both buyers set pre-agreed prices for each demand state in all periods, denoted by $\mathbf{p}^c = (p_H^c, p_L^c)$, which can depend on transparency λ.

For any given choice of $\mathbf{p}^c$, buyer and seller welfare are measured using the definitions in Section 4.1. In particular, each buyer's expected profit averaged over demand states is denoted by $\bar{\pi}(\mathbf{p}^c, \lambda)$, and the single-period expected profit under demand state $i \in \{H, L\}$ is denoted by $\pi_i(\mathbf{p}^c, \lambda)$. We use the average expected profit when evaluating welfare.

Given a cooperative strategy $\mathbf{p}^c$, the buyers play a repeated game using a *grim trigger* strategy: initially, both buyers cooperate and price according to $\mathbf{p}^c$, but as soon as one defects by setting a price that differs from $\mathbf{p}^c$, the other will also defect for the remainder of the game, thereby reverting to the non-cooperative Nash equilibrium, $\mathbf{p}^*(\lambda)$. We say that a cooperative strategy $\mathbf{p}^c$ is a *collusive outcome* under transparency level λ if there is at least one buyer each period making higher expected profit under $\mathbf{p}^c$ than under $\mathbf{p}^*(\lambda)$, i.e., if $\pi_i(\mathbf{p}^c, \lambda) > \pi_i(\mathbf{p}^*(\lambda), \lambda)$ for some $i \in \{H, L\}$. If a collusive outcome is sustainable under grim trigger as a subgame-perfect Nash equilibrium, we say that it is a *collusive equilibrium.*

5.2. Strong Pareto Improvement Under Collusion

We start with the most commonly considered cooperative strategy, where the buyers collaborate to maximize their joint expected profit for each period. This jointly profit-maximizing cooperative strategy, denoted by $\mathbf{p}^{c*}(\lambda) = (p_H^{c*}(\lambda), p_L^{c*}(\lambda))$, is a natural choice since it maximizes buyer welfare. Proposition 5 characterizes $\mathbf{p}^{c*}(\lambda)$ for all $\lambda \in (0,1]$.

Proposition 5. *For any transparency level* $\lambda \in (0,1]$*, there exists a unique jointly profit-maximizing cooperative strategy* $\mathbf{p}^{c*}(\lambda)$*, given by* $(p_H^{c*}(\lambda), p_L^{c*}(\lambda)) = (\underline{p} + p_\Delta^{c*}(\lambda), \underline{p})$*. The equilibrium price dispersion* p_Δ^{c*} *is given by* $p_\Delta^{c*}(\lambda) = 0$ *for all* $\lambda \leqslant \frac{t}{(\gamma+h)[2F(\Delta;\sigma)-1]}$*; otherwise,*

$$p_\Delta^{c*}(\lambda) = (\gamma + h)\left[F\left(\Delta - \frac{\lambda}{t} p_\Delta^{c*}(\lambda); \sigma\right) - \frac{1}{2}\right] - \frac{t}{2\lambda} > 0. \tag{12}$$

Under any λ *such that* $\mathbf{p}^{c*}(\lambda) \neq \mathbf{p}^*(\lambda)(\lambda)$*,* $\mathbf{p}^{c*}(\lambda)$ *is a collusive outcome.*

In other words, the low-demand buyer always keeps his price at $\underline{p}$, while the high-demand buyer may increase his price by p_Δ^{c*} to signal higher demand. Parallel to the competitive market condition in Lemma 1, the necessary and sufficient condition for transparency to improve sellers' welfare is that $p_\Delta^{c*}(\lambda) > 0$. Under collusion, the same condition is also necessary and sufficient for buyers to strictly benefit. This leads to Proposition 6.

Proposition 6 (Demand Signaling Condition)**.** *All sellers and buyers are strictly better off when buyers collude to maximize the joint profit if and only if* $\Delta > \underline{\Delta}^c, \gamma + h > t$, *and* $\lambda > \hat{\lambda}^c$. *Herein,* $\underline{\Delta}^c$ *is any solution satisfying* $F(\underline{\Delta}^c; \sigma) = \frac{1}{2} + \frac{t}{2(\gamma+h)}$, *and* $\hat{\lambda}^c = \frac{t}{(\gamma+h)[2F(\Delta;\sigma)-1]}$.

Proposition 6 establishes that two-sided benefits under jointly profit-maximizing collusion also requires three joint conditions.

(1) High contractual imbalance that dominates random supply variations, $\Delta > \underline{\Delta}^c$.

(2) High underage *<u>or</u>* *overage cost,* $\gamma + h > t$.

(3) Sufficiently high price transparency, $\lambda > \hat{\lambda}^c$.

Under price collusion, two-sided benefits arise only if the gains from demand signaling are compelling enough for buyers to differentiate their prices. When this happens, sellers enjoy improved prices while buyers benefit from reduced loss. We thus refer to this set of conditions as the *demand signaling condition*. In particular, the first condition $\Delta > \underline{\Delta}^c$ highlights a similar interplay between supply uncertainty and contractual imbalance in the collusive setting: the required threshold $\underline{\Delta}^c$ decreases with lower σ as demand signaling becomes more effective. Compared to the three similar conditions under the competitive equilibrium (see Theorem 2), the lack of price competition means that higher underage and overage costs now both promote two-sided benefits by making demand signaling more appealing. Moreover, there is no longer the need for an upper limit on transparency to protect buyer profit.

Figure 6a illustrates $p_\Delta^{c*}(\lambda)$ as a function of transparency λ. In comparison with price dispersion under competitive pricing (see Figure 3b), the effect of buyers' demand signaling through price dispersion is mostly preserved under price collusion. Moreover, similar to Figure 3b, price dispersion decreases as uncertainty σ increases. Figure 6b and Figure 6c also illustrates how the expected profit $\bar{\pi}(\mathbf{p}^{c*}, \lambda)$ and expected mismatch $s(\mathbf{p}^{c*}, \lambda)$ change with λ under different values of σ, in parallel with Figure 4a and 4b.

The jointly profit-maximizing strategy $\mathbf{p}^{c*}$ is only one of the many different cooperative strategies that buyers may engage in to avoid price competition and improve profits. To determine which collusive outcomes are more likely to occur, we also need to consider *stability* of collusion.

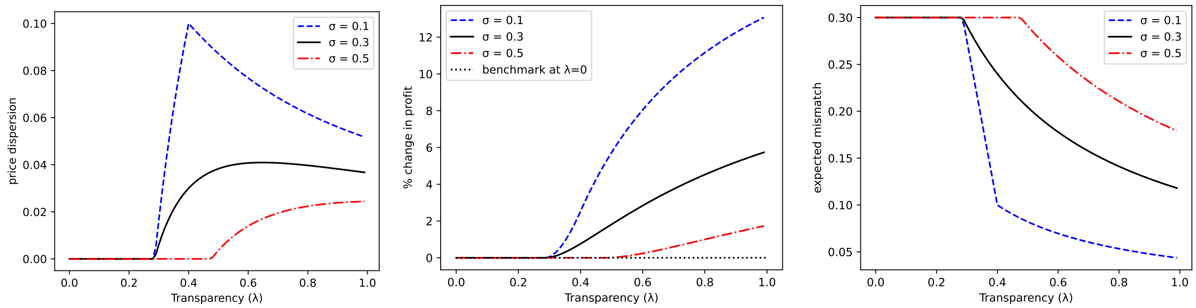

(a) p_Δ^{c*} v.s. λ (b) % Chg in $\bar{\pi}(\mathbf{p}^{c*}, \lambda)$ v.s. λ (c) $s(\mathbf{p}^{c*}, \lambda)$ v.s. λ

Figure 6 Change in buyer's expected profit and the expected supply-demand mismatch as a function of λ, in markets with uniformly distributed supply uncertainty $\epsilon \sim [-\sigma, \sigma]$. The other parameters are $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2, \Delta = 0.3$.

In order for collusion to be sustained in a given period, both buyers need to have no incentive to defect. For a given discount factor δ, this occurs if and only if the discounted value of total expected profit under sustained collusion exceeds that under a defect, i.e., if

$$\pi_i(\mathbf{p}^c, \lambda) + \frac{\delta}{1-\delta}\bar{\pi}(\mathbf{p}^c, \lambda) \geqslant \max_{p_i} \pi_i(p_i, p_j^c, \lambda) + \frac{\delta}{1-\delta}\bar{\pi}(\mathbf{p}^*(\lambda), \lambda) \quad \forall i \in \{L, H\}. \tag{13}$$

The LHS is the discounted total expected profit going forward if no one defects for the remainder of the game, and the RHS is the discounted total expected profit if a buyer defects while in demand state i with optimal deviation, $\max_{p_i} \pi_i(p_i, p_j^c, \lambda)$, and cooperation is discontinued. The minimum value of δ needed for Equation (13) to hold, denoted by $\hat{\delta}(\mathbf{p}^c, \lambda)$, is commonly called the *threshold discount factor* for the given collusive outcome. An outcome can be sustained in collusive equilibrium if and only if the buyers are sufficiently patient, i.e., $\delta \geqslant \hat{\delta}(\mathbf{p}^c, \lambda)$. We therefore say that a collusive equilibrium (or outcome) is *more stable* than another if it requires a *lower* threshold discount factor to sustain.

Theorem 3 establishes that collusive outcomes that give rise to strong Pareto improvement are more stable and profitable than collusion at monopolistic pricing $\underline{\mathbf{p}}$ that leads to no improvement for sellers. The former are therefore more likely to occur as an equilibrium if buyers collude.

Theorem 3. *In any market satisfying both the competitive market and demand signaling conditions, if $\underline{\mathbf{p}}$ is a collusive equilibrium under any given $\lambda > \max\{\hat{\lambda}, \hat{\lambda}^c\}$, then there also exist strongly Pareto-improving collusive equilibria $\mathbf{p}^c$ that are strictly more stable and more profitable than $\underline{\mathbf{p}}$, i.e., $\hat{\delta}(\mathbf{p}^c, \lambda) < \hat{\delta}(\underline{\mathbf{p}}, \lambda)$ and $\bar{\pi}(\mathbf{p}^c, \lambda) > \bar{\pi}(\underline{\mathbf{p}}, \lambda)$. Herein, $\hat{\lambda}^c = \frac{t}{(\gamma+h)[2F(\Delta;\sigma)-1]}$ and $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$.*

The intuition behind greater stability is that efficient price signals make the high-demand buyer less likely to defect, in order to attract more supply and avoid the underage costs. This suggests that even if buyers collude as a result of platform transparency, it is more likely that they will differentiate their prices to signal demand for greater profits. The resulting collusive equilibrium leads to higher market prices that, over time, also benefit all of the sellers.

6. Model Extensions

In this section, we explore three extensions to our model and show how our main result (Theorem 2) applies to more general settings.

6.1. General forms of supply uncertainty

Our base model assumes supply uncertainty is additive and independent of transparency λ and the expected supply y for tractability. In practice, however, two main sources of uncertainty could introduce dependency between σ and either λ or y. First, sellers may consider factors other than prices when deciding whom to trade with. Second, individual supply quantities may vary due to yield uncertainty and harvest decisions. In Appendix A.3, we demonstrate the presence of both factors in our partner platform's market (see Figure 13) and present explicit models that account for these sources of uncertainty. In this section, we extend our main results under the base model to general forms of supply uncertainty that are influenced by either transparency ($\sigma(\lambda)$) or expected supply ($\sigma(y)$), encompassing most of the models discussed in Appendix A.3.

6.1.1. Supply uncertainty varies with transparency, $\sigma(\lambda)$. Creating price transparency may affect how sellers make selling decisions, which in turn may change the degree of supply uncertainty faced by buyers. To capture this, we let supply uncertainty $\epsilon \sim F$ depend on a scale parameter $\sigma(\lambda)$ that is a function of transparency λ, and write $\epsilon \sim F(\cdot;\sigma(\lambda))$.

When transparency creates more supply uncertainty, i.e., when $\sigma(\lambda)$ is an increasing function, transparency-induced uncertainty has two adverse effects on buyers' profit. First, it increases the buyer's expected loss $L(\mathbf{p}^*,\lambda)$ directly due to a more spread out probability distribution of supply variation, which induces greater expected overage or underage. Moreover, it also weakens the effect of demand signaling in equilibrium (Proposition 2), thereby reducing the benefits of price transparency in supply-demand matching. As a result, two-sided benefits are no longer guaranteed under the conditions of Theorem 2, which suggests that a platform needs to be extra careful about any potential transparency-induced supply uncertainties.

Conversely, when transparency helps reduce supply uncertainty, i.e., $\sigma(\lambda)$ is a decreasing function, lower uncertainty not only directly reduces buyers' expected loss from supply variations, but also enhances the effect of demand signaling, creating additional benefits for buyers. Thus, our previous results (Propositions 2, 6 and Theorem 2, 3) still hold, and we include the proofs in Appendix D. In fact, we can obtain stronger results regarding the existence of two-sided benefits in the competitive equilibrium, as stated in Proposition 7.

Proposition 7. *Suppose $\frac{d\sigma(\lambda)}{d\lambda} \leqslant 0$ for all $\lambda \in (0,1]$. Consider any market satisfying the competitive market condition, i.e., $\hat{\lambda} < 1$ and $\lambda > \hat{\lambda}$, where $\hat{\lambda}$ is defined by the implicit function* $(\gamma +$

$h)F(\Delta;\sigma(\hat{\lambda})) - h - \frac{t}{\hat{\lambda}} + 1 - \underline{p} = 0$. Then, if $\frac{d\sigma(\lambda)}{d\lambda} < 0$ for some $\lambda \in (0, \hat{\lambda}]$, there always exists a threshold level $\lambda^* \in (\hat{\lambda}, 1]$, such that any transparency $\lambda \in (\hat{\lambda}, \lambda^*)$ is strongly Pareto improving under the competitive equilibrium.

Proposition 7 establishes that if transparency strictly reduces supply uncertainty, then buyers are guaranteed to benefit under some well-chosen level of λ. Observe that achieving two-sided benefits in this case no longer requires high overage cost h or contractual imbalance Δ. For $\lambda \in (\hat{\lambda}, \lambda^*)$, buyers benefit from reduced expected loss driven by lower supply uncertainty, and sellers benefit from improved price.

6.1.2. Supply uncertainty increases with expected supply, y. When supply uncertainty arises from varying yield from sellers, higher expected supply may lead to greater level of supply uncertainty (see Appendix A.3). To capture this effect, we consider supply uncertainty $\epsilon \sim F(\cdot;\sigma(y))$ where the scale parameter $\sigma(y)$ is a differentiable function non-decreasing in y. We consider general forms of $\sigma(y)$ that are weakly increasing in y and preserve concavity of buyer's profit function:

Assumption 3. *$\sigma(y)$ is weakly increasing in y, and preserves concavity of individual buyers' expected profit, $\pi_i(p_i, p_j, \lambda)$, with respect to the pricing decision, p_i.*

In Appendix A.3, we show graphically that the profit functions under the yield uncertainty models in Appendix A.3 are indeed concave, satisfying Assumption 3.

Since higher expected supply introduces greater uncertainty in this setting, price competition is reduced in the market. Buyers are less inclined to raise prices to secure more supply, because doing so increases uncertainty. This leads to a generalized version of Theorem 2, as stated in Theorem 4.A. In particular, we introduce the term ζ, defined as $\zeta = \int_{-\infty}^{-\Delta} \left[\frac{\partial F(\epsilon;\sigma(y))}{\partial y} \Big|_{y=1} \right] d\epsilon \geqslant 0$, which captures the effect of increasing uncertainty with expected supply, y. Our previous discussions with constant σ can be viewed as a special case where $\zeta = 0$.

Theorem 4.A. *Consider any market satisfying Assumption 3 and the competitive market condition, i.e., $\hat{\lambda} < 1$ and $\lambda > \hat{\lambda}$, where $\sigma_0 = \sigma(1)$, $\zeta = \int_{-\infty}^{-\Delta} \left[\frac{\partial F(\epsilon;\sigma(y))}{\partial y} \Big|_{y=1} \right] d\epsilon \geqslant 0$, and $\hat{\lambda} = \frac{t}{(\gamma+h)[F(\Delta;\sigma_0)-\zeta]-h+1-\underline{p}}$.*

Then, if $\Delta > \underline{\Delta}$, $h \geqslant 1 - \underline{p}$, there always exists a threshold level $\lambda^* \in (\hat{\lambda}, 1]$ such that any $\lambda \in (\hat{\lambda}, \lambda^*)$ is strongly Pareto improving under the competitive equilibrium. Herein, $\underline{\Delta}$ is the smallest value satisfying $F(\underline{\Delta};\sigma_0) = \frac{\gamma+1-\underline{p}}{\gamma+h}$. Moreover, if $\zeta > 0$, such λ^* exists if $\Delta > \underline{\Delta}'$, $h + (\gamma+h)\zeta > 1 - \underline{p}$, where $\underline{\Delta}'$ is the largest value satisfying $F(\underline{\Delta}';\sigma_0) = \frac{\gamma+1-\underline{p}}{\gamma+h} - \zeta$.

Since the term ζ reduces price competition, it has the following effects on the three conditions identified previously under Theorem 2. First, it becomes easier for buyers to receive a net benefit from demand signaling due to reduced price competition. This lowers the minimum threshold required for both contractual balance ($\frac{\partial \underline{\Delta}}{\partial \zeta} \leqslant 0$) and overage cost h. The latter is captured by the inequality $h+(\gamma+h)\zeta \geqslant 1-\underline{p}$, as compared to $h \geqslant 1-\underline{p}$ in Theorem 2. Meanwhile, a higher level transparency $\hat{\lambda}$ is required for buyers to start raising their prices above sellers' reservation price ($\frac{\partial \hat{\lambda}}{\partial \zeta} > 0$). In Appendix A.3, we numerically solve for the equilibrium market prices under $\sigma(y) \propto y$ and $\sigma(y) \propto \sqrt{y}$. We show that aside from lower price competition, these models produce similar market dynamics and qualitative insights compared to those observed under constant σ.

6.2. Buyers turn down unwanted supply

Our base model focuses on an environment where buyers always accept all realized supply for tractable analysis. However, when the overage cost is high ($h > 1-p_i$), it may become profitable for buyers to turn down extra supply above Q_i. Since sellers can only visit one buyer each period due to prohibiting transportation cost, supply that is turned down goes to outside options, such as external buyers, self-consumption or wastage.[5] In this section, we show that our main insight regarding the existence of two-sided benefits (Theorem 2) carries over to a market where overage costs are sufficiently high, such that buyers always turn down any supply above contractual demand.

Recall that the stochastic supply function of buyer i is of the form of $y(p_i, p_j, \lambda) + \epsilon_i$. When any extra supply is turned down, the buyer's realized profit becomes

$$\min\{y(p_i, p_j, \lambda) + \epsilon_i, Q_i\} \cdot (1-p_i) - \gamma \cdot [y(p_i, p_j, \lambda) + \epsilon_i - Q_i]^-.$$

Observe that, substituting $s_i(p_i, p_j, \lambda) = y(p_i, p_j, \lambda) - Q_i$, we can equivalently write buyer profit as

$$[y(p_i, p_j, \lambda) + \epsilon](1-p_i) - \gamma[s_i(p_i, p_j, \lambda) + \epsilon]^- - (1-p_i)[s_i(p_i, p_j, \lambda) + \epsilon]^+.$$

In comparison with the original realized profit in Equation (2), the new profit is mathematically equivalent to incurring a *price-dependent* "overage cost" of $h := 1-p_i$. The equilibrium prices no longer admit a closed-form characterization due to additional complexity from i) "overage cost" being a function of price, and ii) asymmetric overage cost incurred by buyers when $p_i \neq p_j$.

For this setting, we present an alternative version of the generalized main result (Theorem 4.A in Section 6.1) in Theorem 4.B below, where buyers reject any supply exceeding their contractual demand, Q_i. We assume that overtime, rejected supply is equally distributed among all sellers. Thus, all sellers are better off when there is price improvement and reduction in the expected amount of rejected supply. In addition, like Theorem 4.A, Theorem 4.B also considers general supply-dependent uncertainty, where the scale parameter $\sigma(y)$ satisfying Assumption 3 with the corresponding parameter, $\zeta \geqslant 0$.

[5] Here, it is assumed that all rejected supply goes to outside options. Our high-level insights remain valid if part of rejected supply is sold to the other buyer in the market. We provide additional analysis for the latter scenario in Appendix A.2.

Theorem 4.B. *Consider any market satisfying Assumption 3 and the competitive market condition, i.e.,* $\hat{\lambda} < 1$ *and* $\lambda > \hat{\lambda}$*, where* $\sigma_0 = \sigma(1)$*,* $\zeta = \int_{-\infty}^{-\Delta} \left[\frac{\partial F(\epsilon;\sigma(y))}{\partial y}\Big|_{y=1} \right] d\epsilon$*, and* $\hat{\lambda} = \frac{t(1-\int_{-\infty}^{-\Delta} F(\epsilon;\sigma_0)d\epsilon)}{(\gamma+1-\underline{p})[F(\Delta;\sigma_0)-\zeta]}$*. Suppose both buyers turn down realized supply above demand* Q_i*.*

Then, if $\Delta > \underline{\Delta}$*, there exists a threshold level* $\lambda^* \in (\hat{\lambda}, 1]$ *such that any* $\lambda \in (\hat{\lambda}, \lambda^*)$ *is strongly Pareto improving under the competitive equilibrium. Herein,* $\underline{\Delta}$ *is the smallest value satisfying* $F(\underline{\Delta};\sigma_0) = 1$*. Moreover, if* $\zeta > 0$*, such a* λ^* *exists if* $\Delta > \underline{\Delta}'$*, where* $\underline{\Delta}'$ *is the largest value satisfying* $F(\underline{\Delta}';\sigma_0) = 1 - \zeta$*.*

Comparing Theorem 4.B and 4.A, the requirements for contractual imbalance and overage costs remain similar, though they no longer depend on the overage cost, h. Specifically, achieving two-sided benefits requires a high Δ or low σ_0 ($\Delta > \underline{\Delta}$), along with sufficiently high overage costs such that buyers are motivated to turn down supply (e.g., a sufficient condition is $h \geqslant 1 - \underline{p}$). In particular, when $h = 1 - \underline{p}$, the required threshold for Δ is identical for when buyers commit to accepting all supply and when buyers turn down additional supply. In the latter case, however, there is slightly higher price competition in the market, reflected in a lower threshold $\hat{\lambda}$ needed for price improvement. Since unwanted supply is consistently declined, buyers are less concerned about potential overages. This leads to increased price competition as buyers aim to secure enough supply to avoid shortages without the risk of over-procurement.

In Appendix A.2, we numerically solve for the equilibrium market prices and demonstrate that other market dynamics in this alternative model also closely resemble those in our base model.

Lastly, it is worth noting that when $\zeta = 0$, the condition $\Delta > \underline{\Delta}$ where $F(\underline{\Delta};\sigma_0) = 1$ may seem stringent for distributions with unbounded support (e.g. normal distribution). While this condition is *sufficient* for two-sided benefits to occur, it is not necessary. We show in Appendix A.2 that under normal distribution, two-sided benefits can still emerge as long as $F(\underline{\Delta};\sigma_0)$ is sufficiently large (see Figure 11d).

7. Managerial Insights

Our analytical results shed light on key factors that drive two-sided benefits under price transparency. This helps information platforms answer two important managerial questions. First, how should they identify target markets where transparency is likely to create two-sided benefits? Second, should they provide full or partial transparency? In this section, we provide managerial guidelines and practical suggestions for how an information platform can answer these questions and implement the appropriate solutions.

7.1. General Recommendations

Target markets. For information providers whose objective is to create two-sided benefits, our results show that it is desirable for their target markets to have two distinct characteristics (Theorems 2, 4.A, 4.B):

1. High overage costs (h) that make excessive supply undesirable.
2. High contractual imbalance that dominates random supply variations (i.e., large $F(\Delta;\sigma)$).

In these markets, improved transparency creates moderate price competition that benefits the sellers and effective demand signaling that benefits the buyers. Notably, whether a specific contractual imbalance is deemed sufficient depends on natural supply variations in the market. A particular imbalance value, Δ, might be high enough to produce two-sided benefits in a market with low supply uncertainty (σ), but insufficient in one with high σ. In the latter case, as Corollary 2.1 suggests, two-sided benefits can still emerge if the platform designer helps create a low-uncertainty environment that allows buyers to signal demand efficiently.

Optimal transparency level. The level of transparency needs to be sufficiently high to induce a price increase by buyers to benefit the sellers; meanwhile, depending on the intensity of price competition in the market, the platform may also need to limit price transparency to an intermediate level to protect buyers' profit (Theorems 2, 4.A, 4.B, Proposition 7). The latter is no longer necessary if the platform expects collusive behavior among buyers (Proposition 6, Theorem 3). Nonetheless, moderate transparency could still be beneficial, as higher transparency could encourage buyer collusion and reduce the welfare benefits for sellers.

When evaluating whether partial or full transparency is desired, the platform designer should consider two factors impacting the intensity of price competition (Proposition 4):

1. High underage costs (γ) increases price competition, and vice versa.
2. High market differentiation (t) reduces price competition, and vice versa.

With high underage costs and little differentiation, price competition arises even with limited transparency (small $\hat{\lambda}$, $\lambda^* << 1$). Meanwhile, excessive transparency can drive prices up and harm the buyers. We call these *high competition* markets, where partial transparency is preferred to protect buyer profits. Conversely, when overage costs are low and market differentiation is high, price competition remains limited even with near-full transparency (large $\hat{\lambda}$, $\lambda^* \approx 1$). In these markets, full transparency is beneficial for both inducing price competition and enabling buyers to signal demand. We refer to these as *low competition* markets.

Table 1 summarizes the above recommendations under different market conditions.

The horizontal axis in the table captures factors conducive for two-sided benefits in target markets. The vertical axis captures key determinants for establishing the optimal level of transparency. Overall, Table 1 highlights that achieving two-sided benefit in target markets requires implementing a well-chosen level of price transparency, and, in some cases, creating a low uncertainty environment for buyers. A platform's optimal strategy depends on the interplay of contractual imbalance, supply uncertainty, underage and overage costs and degree of differentiation in the particular market. We discuss potential approaches and practical challenges of implementing partial transparency and reducing supply uncertainty through platform implementation in Section 7.3.

	High contractual imbalance, costly overage ($\Delta > \underline{\Delta}, h \geqslant 1-\underline{p}$)	**Low contractual imbalance, costly overage** ($\Delta \leqslant \underline{\Delta}, h \geqslant 1-\underline{p}$)
High competition *(High γ, low t, $\lambda^* << 1$)*	Partial information	Partial information & reduce σ
Low competition *(Low γ, high t, $\lambda^* \approx 1$)*	Full information	Full information & reduce σ

Table 1 Recommendations for information platforms to ensure two-sided benefits: what level of information transparency to provide (full or partial) and whether creating a low-uncertainty environment is necessary, depending on market characteristics (Δ, h, γ, t). Herein, $\underline{\Delta}$ is the smallest value satisfying $F(\underline{\Delta};\sigma) = \frac{\gamma+1-\underline{p}}{\gamma+h}$, and λ^* is the maximum value of strongly Pareto improving λ as defined in Proposition 4 .

7.2. Implications for PemPem

PemPem operates in the first mile of a commodity supply chain, a setting with generally high market-wide competition (as also illustrated by the relatively high price sensitivity in Figure 1b), and local monopolies or oligopolies in certain geographies. This suggests that whether PemPem falls in the top or bottom row of Table 1 depends on the geographical region (see also Figure 7). The level of contractual imbalance in the market, on the other hand, may vary over time. While PemPem does not collect detailed data on target contractual quantities of buyers, anecdotally, buyers aim to procure proportional to the contractual deadline. This means buyers' target quantities can be more differentiated closer to deadlines, and less differentiated when far from deadlines. Hence, the appropriate level of transparency and the need for uncertainty reduction may change across different regions or even time periods.

We hence calibrate our model for PemPem by considering reasonable values of the parameters $\underline{p}, t, h, \gamma, \Delta, F$ for multiple scenarios; more details about the estimation process are in Appendix A.1. In general, t can be thought of as the stickiness of informed supply, and can reflect factors such as geographical differentiation, trust or loyalty. Using estimates of supply elasticity based on transaction data in different regional markets, we find high and low competition correspond to $t \approx 0.13$ and $t \approx 0.33$ respectively. The contractual imbalance, Δ, can be thought of as the percentage difference between different buyers' contractual quantities. Back-engineering Δ from price dispersion data yields values ranging from close to 0 to 30%, and anecdotal discussions suggest similar estimates. Thus, we consider $\Delta = 0.3$ during periods of high demand asymmetry and $\Delta = 0.15$ during periods of lower demand asymmetry. Supply uncertainty, F, is calibrated by fitting a normal distribution on bootstrapped supply data. The resultant F follows a normal distribution of mean 0 and standard deviation, $\sigma = 0.21$. Lastly, h, γ are difficult to estimate based on PemPem data, and so we consider cases where $h = \gamma = 1 - \underline{p}$. Results for different values of γ are qualitatively similar, and can be found in Appendix A.1.

Figure 7 illustrates the recommended levels of price transparency and reductions in supply uncertainty for PemPem to achieve two-sided benefits from price transparency, depending on the geographical region where PemPem operates (influencing t) and the level of contractual imbalance in the market at a given point in time (influencing Δ). The structure of the Figure follows the structure of Table 1. Figure 7a and 7b reflect the first row of Table 1, capturing a geographical region with high competition (grey shading). Figure 7c and 7d correspond to the second row of Table 1, showing a geographical region with lower competition (blue shading). Furthermore, the right figure (7a and 7c) of each set captures settings close to contractual deadlines (high Δ), reflecting the first column of Table 1. The other figures on the left (7b and 7d) capture settings further away from contractual deadlines (low Δ). Shaded regions indicate levels of transparency required for Pareto improvement under different levels of supply uncertainty in the market. The red vertical line indicates the estimated level of actual supply uncertainty.

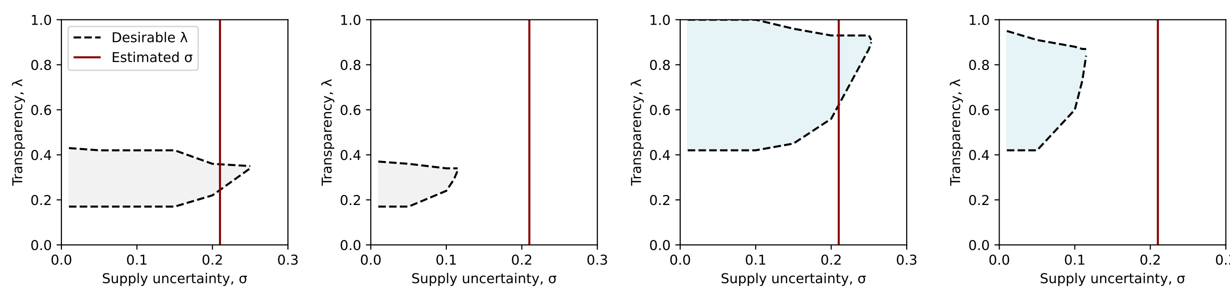

(a) $t = 0.13, \Delta = 0.3$ (b) $t = 0.13, \Delta = 0.15$ (c) $t = 0.33, \Delta = 0.3$ (d) $t = 0.33, \Delta = 0.15$

Figure 7 **Recommended levels of price transparency versus reductions in supply uncertainty for PemPem to achieve two-sided benefits from price transparency, depending on the geographical region where PemPem operates and the level of contractual imbalance in the market at a given point in time.**

In line with Table 1, when t is low (Figure 7a, 7b), the ideal range of λ is around $[0.2, 0.4]$, and so partial transparency is required. When t is high (Figure 7c, 7d), the regions of improvement is much closer to $\lambda = 1$, showing that near-full transparency is desirable. Given appropriate levels of transparency, buyers are likely to benefit from demand signaling during periods of high contractual imbalance (Figure 7a, 7c), but less likely during periods of lower contractual imbalance (Figure 7b, 7d), unless supply uncertainty is reduced to $\sigma \approx 0.1$ or lower.

The above analysis shows that the platform's optimal information disclosure policies needs to be region-sensitive, based on local supply elasticity estimates (i.e., local values of t). This might explain the mixed response of buyers in *PemPem* regarding their willingness to share information, per Section 1. In addition, creating a low-uncertainty environment can encourage more active buyer participation even when contractual imbalances are small. In the next section, we explore how PemPem and other platforms can implement both aspects of platform design.

7.3. Implementing Partial Transparency and Reducing Supply Uncertainty

In this section, we discuss potential pathways and challenges for implementing partial transparency and reducing supply uncertainty in practice. While identifying the most effective methods is beyond the scope of this work, we hope this discussion can motivate further research to address these practical challenges and explore possible solutions.

Implementing Partial Transparency. When partial transparency is necessary, the challenge lies in how to implement it effectively. To ensure fairness among vulnerable sellers, the platform should avoid limiting transparency to a select group of sellers. We show in Appendix A.6 that such restrictions could result in market segmentation and allow buyers to exploit uninformed sellers, exacerbating market inequality. Instead, the platform should aim to provide equal access to all sellers, while limiting disclosed information uniformly across sellers overtime. This approach has been successful in other contexts, such as salary information platforms (e.g., Glassdoor), where only average salaries are disclosed rather than detailed salary distributions to all users. For platforms like PemPem, a recommended strategy for partial transparency is to provide sellers with information from a selected subset of buyers. The amount of shared buyer data could vary based on regional competition and seller needs. Alternatively, the platform could randomize the buyer information available to suppliers or offer aggregated data, such as historical averages or average prices within small buyer groups. Lastly, the optimal transparency level can be challenging to estimate sometimes and may change over time. This will require active re-balancing of the platform's information provision strategy, based on ongoing monitoring of market prices and feedback from both sides of the market.

Reducing supply uncertainty. A similar question arises of how supply uncertainty can be reduced in practice. In one of our uncertainty models, we show that increased transparency can help reducing uncertainty if it also reduces the noise in seller decision making (see model 2 in Appendix A.3). The platform could achieve this goal by helping informed sellers make rational selling decisions. For instance, informed suppliers' selling decisions can be influenced by irregular harvest schedules, or implicit costs of switching trading partners, such as transaction costs, trust, and other psychological, non-economic costs unknown to the buyers (see, e.g., Klemperer (1995)). To reduce randomness in such decisions, designs such as secured digital payment systems and buyer certification can be introduced to minimize switching costs and facilitate rational selling decisions based on price. Alternatively, the platform may also explore centralized supply-demand matching in order to minimize the risk of shortages or surpluses – an approach that *PemPem* is now indeed pursuing. The results in this paper provide strong theoretical support that design choices of this kind can play a central role in improving efficiency and achieving two-sided benefits.

Reducing reliance on buyer participation. Both partial information and uncertainty reduction aim to protect buyer profit and hence encourage buyer participation. While crucial for two-sided benefits, these strategies can be impractical in some cases, such as when sellers resist receiving only partial information or when yield uncertainty is unavoidable. When this happens, a platform may shift its focus to addressing potential buyer reluctance more proactively. For instance, platforms that rely heavily on buyer-shared information like *Pempem* could offer additional financial incentives to crowd-source data from sellers, especially in high-competition regions, thereby reducing the platform's dependence on buyer participation in the long run.

8. Concluding Remarks

Motivated by smallholder supply chains in developing markets, this paper examines whether and how increased price transparency creates two-sided benefits in a setting with costly underage and overage, contractual imbalance and geographically differentiated products. Our analyses reveal important nuances about the interplay between transparency, pricing, contractual obligations in the supply chain, and supply uncertainty. In doing so, we close several gaps between the empirical literature and the theoretical economic literature on this topic, and offer possible explanations for the variations in empirical findings.

By fully characterizing the competitive equilibrium, we uncover two mechanisms that drive welfare improvements in the market: price competition favoring the sellers and demand signaling benefiting the buyers. We use this to identify the following three conditions for two-sided benefits. *(1) High contractual imbalance with low supply uncertainty:* Buyers facing greater contractual imbalance gain more from demand signaling as it reduces supply-demand mismatch. Low supply uncertainty further amplifies such benefits of demand signals. *(2) High overage costs:* A high overage cost makes additional supply less desirable. This benefits buyers by keeping price competition low and increasing the benefit of demand signaling. *(3) Well-chosen levels of transparency:* A minimum transparency level ensures price competition that benefits sellers, while a maximum transparency limit may help protect buyers. Thus, partial transparency can be desirable in high-competition markets with low differentiation and high underage penalties, whereas full transparency is ideal in low-competition markets.

The analysis of the repeated game further shows that these benefits can be preserved under price collusion. The insights in the paper equivalently hold, with a minimal change of variables needed, in a model where buyers and sellers are reversed, i.e., the sellers play the role of price-setting newsvendors and the buyers search for price.

Based on our findings, we provide managerial recommendations for platform designers to identify target markets and the desirable transparency level to implement. We also calibrate this approach to the context of our collaborating platform to guide platform design.

Our analysis reveals the benefits of ensuring the appropriate transparency level and minimizing uncertainty in information platform design. However, it does not fully address the most effective methods for implementing these strategies in practice. As discussed in Section 7.3, successful implementation demands careful consideration of factors like fairness, user behavior, and local market conditions. We hope this work inspires further research to offer more practical guidance for platform designers and social planners in achieving these design goals.

References

Abraham R (2006) Mobile phones and economic development: Evidence from the fishing industry in india. *2006 International Conference on Information and Communication Technologies and Development*, 48–56 (IEEE).

Aker JC (2008) Does digital divide or provide? the impact of cell phones on grain markets in niger. *Center for Global Development working paper* (154).

Aker JC (2010) Information from markets near and far: Mobile phones and agricultural markets in niger. *American Economic Journal: Applied Economics* 2(3):46–59.

Aker JC, Fafchamps M (2015) Mobile phone coverage and producer markets: Evidence from west africa. *The World Bank Economic Review* 29(2):262–292.

Baye MR, Morgan J (2001) Information gatekeepers on the internet and the competitiveness of homogeneous product markets. *American Economic Review* 91(3):454–474.

Bergemann D, Morris S (2019) Information design: A unified perspective. *Journal of Economic Literature* 57(1):44–95.

Borodin V, Bourtembourg J, Hnaien F, Labadie N (2016) Handling uncertainty in agricultural supply chain management: A state of the art. *European Journal of Operational Research* 254(2):348–359.

Burdett K, Judd KL (1983) Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society* 955–969.

Chen YJ, Shanthikumar JG, Shen ZJM (2015) Incentive for peer-to-peer knowledge sharing among farmers in developing economies. *Production and Operations Management* 24(9):1430–1440.

Chen YJ, Tang CS (2015) The economic value of market information for farmers in developing economies. *Production and Operations Management* 24(9):1441–1452.

Chen Z, Liu F (2021) Multi-outsourcing supply chain coordination under yield and demand uncertainties. *Expert Systems with Applications* 181:115177.

Ellison G, Ellison SF (2009) Search, obfuscation, and price elasticities on the internet. *Econometrica* 77(2):427–452.

Fafchamps M, Minten B (2012) Impact of sms-based agricultural information on indian farmers. *The World Bank Economic Review* 26(3):383–414.

Futch MD, McIntosh CT (2009) Tracking the introduction of the village phone product in rwanda. *Information Technologies & International Development* 5(3):pp–54.

Gokarn S, Kuthambalayan TS (2017) Analysis of challenges inhibiting the reduction of waste in food supply chain. *Journal of cleaner production* 168:595–604.

Goyal A (2010) Information, direct access to farmers, and rural market performance in central india. *American Economic Journal: Applied Economics* 2(3):22–45.

Green EJ, Porter RH (1984) Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society* 87–100.

Güler K, Körpeoğlu E, Şen A (2018) Newsvendor competition under asymmetric cost information. *European Journal of Operational Research* 271(2):561–576.

He QC, Chen YJ, Shen ZJ (2018) On the formation of producers' information-sharing coalitions. *Production and Operations Management* 27(5):917–927.

Jensen R (2007) The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The quarterly journal of economics* 122(3):879–924.

Keren B (2009) The single-period inventory problem: extension to random yield from the perspective of the supply chain. *Omega* 37(4):801–810.

Klemperer P (1995) Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The review of economic studies* 62(4):515–539.

Knapp S, van der Heijden MG (2018) A global meta-analysis of yield stability in organic and conservation agriculture. *Nature communications* 9(1):1–9.

Kremer I, Mansour Y, Perry M (2014) Implementing the "wisdom of the crowd". *Journal of Political Economy* 122(5):988–1012.

Levi R, Rajan M, Singhvi S, Zheng Y (2020) The impact of unifying agricultural wholesale markets on prices and farmers' profitability. *Proceedings of the National Academy of Sciences* 117(5):2366–2371.

Liao CN, Chen YJ, Tang CS (2019) Information provision policies for improving farmer welfare in developing countries: Heterogeneous farmers and market selection. *Manufacturing & Service Operations Management* 21(2):254–270.

Lippman SA, McCardle KF (1997) The competitive newsboy. *Operations research* 45(1):54–65.

Newman N, Bergquist LF, Immorlica N, Leyton-Brown K, Lucier B, McIntosh C, Quinn J, Ssekibuule R (2018) Designing and evolving an electronic agricultural marketplace in uganda. *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, 1–11.

Palley AB, Soll JB (2019) Extracting the wisdom of crowds when information is shared. *Management Science* 65(5):2291–2309.

Papanastasiou Y, Bimpikis K, Savva N (2018) Crowdsourcing exploration. *Management Science* 64(4):1727–1746.

Parker C, Ramdas K, Savva N (2016) Is it enough? evidence from a natural experiment in india's agriculture markets. *Management Science* 62(9):2481–2503.

Parlar M (1988) Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Research Logistics (NRL)* 35(3):397–409.

Petruzzi NC, Dada M (1999) Pricing and the newsvendor problem: A review with extensions. *Operations research* 47(2):183–194.

Rasch A, Herre J (2013) Customer-side transparency, elastic demand, and tacit collusion under differentiation. *Information Economics and Policy* 25(1):51–59.

Ritchie H (2021) Smallholders produce one-third of the world's food, less than half of what many headlines claim. *Our World in Data* .

Salop S, Stiglitz J (1977) Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies* 44(3):493–510.

Schultz C (2005) Transparency on the consumer side and tacit collusion. *European Economic Review* 49(2):279–297.

Schultz C (2017) Collusion in markets with imperfect price information on both sides. *Review of Industrial Organization* 50(3):287–301.

Sodhi MS, Tang CS (2014) Supply-chain research opportunities with the poor as suppliers or distributors in developing countries. *Production and operations management* 23(9):1483–1494.

Stahl DO (1989) Oligopolistic pricing with sequential consumer search. *The American Economic Review* 700–712.

Stigler GJ (1964) A theory of oligopoly. *Journal of political Economy* 72(1):44–61.

Svensson J, Yanagizawa D (2009) Getting prices right: the impact of the market information service in uganda. *Journal of the European Economic Association* 7(2-3):435–445.

Tsay AA, Agrawal N (2000) Channel dynamics under price and service competition. *Manufacturing & Service Operations Management* 2(4):372–391.

Varian HR (1980) A model of sales. *The American economic review* 70(4):651–659.

Xiao S, Chen YJ, Tang CS (2020) Knowledge sharing and learning among smallholders in developing economies: Implications, incentives, and reward mechanisms. *Operations Research* 68(2):435–452.

Zhou J, Fan X, Chen YJ, Tang CS (2021) Information provision and farmer welfare in developing economies. *Manufacturing & Service Operations Management* 23(1):230–245.

Appendix A: Supplemental Material

A.1. Model Calibration

As an information platform, Pempem collects transaction data shared by both sellers and buyers in the market. In our setting, sellers are middlemen who aggregate palm fruits from farmers and sell to the buyers, mostly loading ramp owners. One typical entry of transaction data includes unique identifiers of the seller and the buyer, price per kilogram of palm fruits, and the quantity of palm fruits.

Estimating t. To estimate t in the case of PemPem, for each buyer i, we regress the fraction y_i of informed sellers who sell to i, on the difference in prices between buyer i and competing buyers j,

$$y_i = \beta \cdot (p_i - p_j) + c. \tag{14}$$

Since the fraction of informed sellers who sell to i in our duopoly setting is given by $y_i = \frac{1}{2t}(p_i - p_j) + \frac{1}{2}$ if $|p_i - p_j| \leqslant t$ (see second part of Equation (1)), we have $\beta = \frac{1}{2t}$. We thus estimate the implied value of t using $\hat{t} = \frac{1}{2\beta}$. Note that $c = \frac{1}{2}$ in our two-buyer model, which represents the fraction of informed sellers captured without price differentiation. Here, we fit both β and c for each buyer, as c could vary in practice. We estimate the value of y_i and $p_i - p_j$ each week for every chosen buyer during the time period from 2021-11-02 to 2022-06-29. This yields 28 to 33 data points for various buyers depending on data availability.

- **Estimating $p_i - p_j$.** Since Pempem operates in a multi-buyer market, we interpret $p_i - p_j$ as buyer i's average weekly price, p_i, relative to the average weekly price averaged over all his competing buyers, p_j. In other words, we aggregate all competing buyers as one buyer j in order to apply our two-buyer model. Note that each buyer may only compete with a subset of all buyers on the platform, thus we further define 'competing buyers' as the subset of buyers who have common sellers as buyer i. In the model, all prices are normalized by the downstream per unit revenue; we thus normalize the price difference by the average weekly price of i and j, assuming the buyers operate with a narrow profit margin so that the per-unit market price is a good approximation for downstream per unit revenue.
- **Estimating y_i.** In Equation (14), y_i measures the fraction of informed sellers who have options to sell to both i and j and choose to sell to i. We estimate y_i using

$$y_i = \frac{\text{Weekly total \# of seller-shared transactions with buyer } i}{\text{Weekly total \# of seller-shared transactions with buyer } i \text{ and competing buyers } j}.$$

- **Buyer selection criteria.** We choose buyers on the platform as buyer i for analysis based on the following criteria: 1) the total number of seller-shared prices from 2021-11-02 to 2022-06-29 is at least 200; 2) there are some variety in the price-sharing sellers, such that less than 90% of the seller-shared prices about the buyer are shared by one single seller. 3) c is not too small ($c \geqslant 0.2$), such that the setting is closer to a two-buyer model with $c = \frac{1}{2}$ in the duopoly setting. Four buyers in total satisfy the three criteria, and we report all regression results for the four buyers.

Table 2 summarizes the estimated values of t for four selected buyers. Figure 1b in Section 1 illustrates the plot of y_i versus $p_i - p_j$ for Buyer 2. In our analysis, we thus chose $t = 0.33$ and $t = 0.13$ for low and high competition scenarios.

	Slope	R^2	Implied $\hat{t}$
Buyer 1	1.26 (1.65)	0.02	0.40
Buyer 2	1.63 (0.48)***	0.30	0.33
Buyer 3	3.84 (1.73)**	0.18	0.13
Buyer 4	1.54 (0.71)**	0.17	0.32

Table 2 Estimates for t for the 4 buyers with the most sales data on PemPem. **, and * denote (two-tailed) significance at the 5% and 1% respectively, for regression coefficient estimates. Standard errors are given in parentheses.**

Estimating supply uncertainty. When data collected on the platform is complete, supply uncertainty faced by the buyers can be directly estimated using the empirical distribution of supply acquired by individual buyers. In the case of Pempem, collection of sales data relies on self-reporting, and so the empirical distribution of supply may not reflect the true distribution due to missing data.

We thus select one seller who most consistently reports his sale quantity on a daily basis, and obtain the empirical distribution of weekly supply by individual sellers based on quantity data shared by the chosen seller. Both anecdotal discussions and data suggest each buyer receives supply from around 10 different sellers every week. Empirical distribution of each buyer's weekly supply is hence estimated using the sampled distribution of repeatedly sampling and summing 10 independent samples from the distribution of individual seller supply. The resultant empirical distribution is scaled to mean = 1 (since each buyer's expected supply is normalized to 1 in our model when there is no price dispersion) and then shifted to mean = 0 (i.e., approximate normalized empirical supply uncertainty $\tilde{X} = \frac{X}{\text{mean(X)}} - 1$). A normal distribution of mean 0 and $\text{sd}(\tilde{X})$ is fitted for numerical calculations (see Figure 8). This yields $\sigma = \text{sd}(\tilde{X}) = 0.21$ and estimated supply uncertainty $F(\cdot,\sigma) \sim \mathcal{N}(0, 0.21)$.

Estimating Δ. We estimate the value of Δ based on the level of price dispersion in the market. Recall that at sufficiently high transparency, the equilibrium price dispersion is given by

$$p^*_\Delta(\lambda) = \frac{\gamma + h}{3}\left[1 - 2F\left(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta; \sigma\right)\right].$$

We use $\lambda = 1$, $\gamma = h = 0.4$, and $F(\cdot,\sigma) \sim \mathcal{N}(0, 0.21)$ as calculated above.

For values of t and p^*_Δ, we use the data for the four buyers in Table 2 and their competing buyers since data for these buyers are the most complete. For instance, for the case of Buyer 2, we use $t = 0.33$, and obtain all prices of Buyer 2 and his competing buyers on day d, denoted as $\{p_i\}_d$, for all d during the time period from 2021-11-02 to 2022-06-29. Similarly, all values of p_i are normalized by the daily market price averaged across all buyers. The daily price dispersion in Buyer 2's local market is calculated using $p^*_\Delta(d) = \max_i\{p_i\}_d - \min_i\{p_i\}_d$. To estimate upper and lower bound of Δ, we use the maximum and minimum daily price dispersion during the time period, i.e., $\max_d p^*_\Delta(d)$ and $\min_d p^*_\Delta(d)$. All four buyers' markets give similar estimates where $\max_d p^*_\Delta(d) \approx 0.07 \sim 0.08$ and $\min_d p^*_\Delta(d) \approx 0.01$.

Putting all parameter estimates together, we solve the above equation for maximum Δ and minimum Δ using maximum and minimum daily price dispersion. This yields the estimation that Δ ranges from 0.04 to 0.3 in the market. This estimate is consistent with anecdotal evidence from the field team.

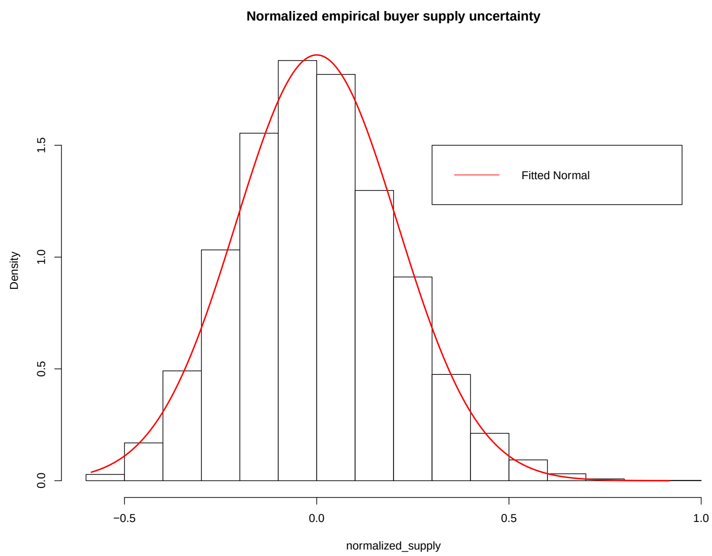

Figure 8 Approximate supply uncertainty from empirical supply distribution

Estimating $\underline{p}$. The final parameter to estimate is the reservation price, $\underline{p}$. In our model, $\underline{p}$ is interpreted as a lower bound for monopolistic pricing, normalized by the downstream per unit revenue. In other words, for each unit of supply, $1-\underline{p}$ denotes the maximum profit margin a buyer earns as a percentage of the converted downstream revenue.

So far, we have focused on the primary palm fruit market where sellers are middlemen and buyers are loading ramps. In addition to this market, Pempem data also include some sales where loading ramps sell their fruits to to local mills at a higher price. The market between loading ramps and mills can thus be treated as the 'downstream' market of the primary market between middlemen and loading ramps. We estimate $\underline{p}$ using the minimum price ratio between the primary market and the downstream market, i.e.,

$$\underline{p} = \min_d \quad \frac{\text{minimum primary market price on day } d}{\text{maximum downstream price } d}.$$

This yields $\underline{p} \approx 0.6$, in line with anecdotal evidence in the field.

Figures 9 and 10 show the recommended transparency levels under different values of γ (while keeping other parameters the same as Figure 7).

A.2. Additional Numerical Results for Section 6.2

Figure 11 demonstrate market prices, the expected quantity of waste (i.e., rejected supply) and buyer profit when excess supply is turned down by the buyers. The market dynamics closely resemble those under the base model. In particular, the blue dotted line and the black solid line in Figure 11d demonstrate that buyers can still benefit (and therefore leading to two-sided benefits) under normally distributed uncertainty with sufficiently small variance, even though the sufficient condition ($\Delta > \underline{\Delta}$ and $F(\underline{\Delta};\sigma_0) = 1$) in Theorem 4.B is not satisfied.

We provide a brief discussion on the alternative scenario where part of rejected supply is sold to the other buyer in the market, and show that our high-level insights of Section 6.2 remain valid.

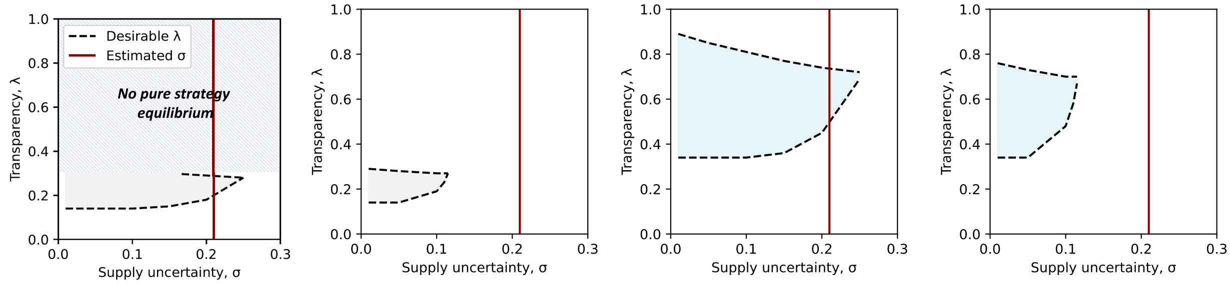

(a) $t = 0.13, \Delta = 0.3$ (b) $t = 0.13, \Delta = 0.15$ (c) $t = 0.33, \Delta = 0.3$ (d) $t = 0.33, \Delta = 0.15$

Figure 9 $\gamma = 0.6$

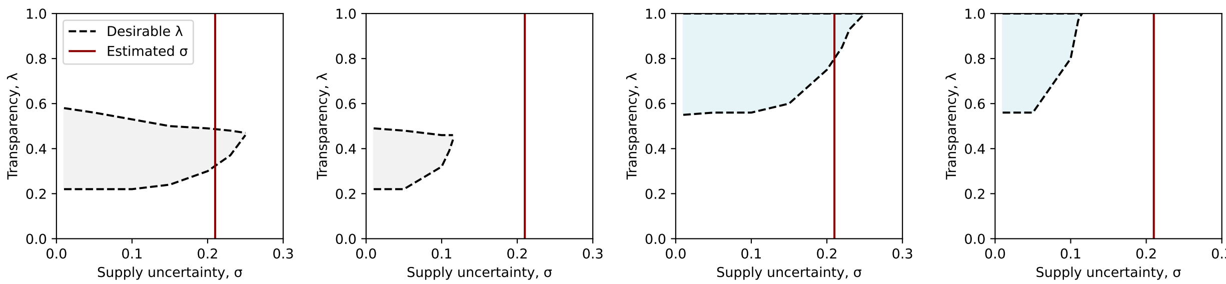

(a) $t = 0.13, \Delta = 0.3$ (b) $t = 0.13, \Delta = 0.15$ (c) $t = 0.33, \Delta = 0.3$ (d) $t = 0.33, \Delta = 0.15$

Figure 10 $\gamma = 0.2$

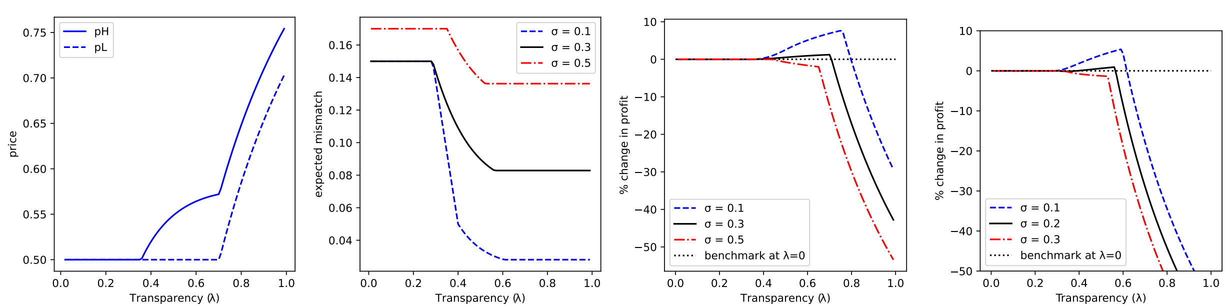

(a) $p^*(\lambda)$ for $\sigma = 0.3$ (b) waste (c) profit (d) profit (normal)

Figure 11 **Buyers turn down additional supply. (a)-(c) shows market dynamics under uniform supply uncertainty $\epsilon \sim U[-\sigma, \sigma]$, and (d) shows buyer profit under normal supply uncertainty, with $\epsilon \sim \mathcal{N}(0, \sigma^2)$. The other parameters are $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.25, \Delta = 0.3$.**

Consider two scenarios with different correlation between supply uncertainty faced by either buyers, ϵ_i, ϵ_j. First, when ϵ_i, ϵ_j are positively correlated, re-direction of supply to the other buyer rarely happens, since it is likely that both buyers either face supply surplus or supply shortage during the same period. Intuitively, re-direction of supply does not significantly alter market dynamics due to its rare occurrence.

Second, when ϵ_i, ϵ_j are negative correlated, re-direction of supply to the other buyer becomes more common. We will consider the extreme case where $\epsilon_i = -\epsilon_j$, such that total supply in the market is fixed. For simplicity, we focus on the case where $Q_i + Q_j = 2$ such that total supply always matches total demand. Whenever a buyer turns down s units of surplus, the other buyer faces s units of shortage. Suppose an $\alpha \in (0,1)$ fraction of s is now sold to the buyer facing shortage. In this setting, for any realized supply uncertainty ϵ_i, the buyers' realized profit is $Q_i(1-p_i)$ if $\epsilon_i \geqslant Q_i - y(p_i,p_j,\lambda)$, and

$$\underbrace{[y(p_i,p_j,\lambda)+\epsilon_i](1-p_i)}_{\text{revenue from original supply}} + \underbrace{\alpha[Q_i - y(p_i,p_j,\lambda)-\epsilon_i](1-p_i)}_{\text{revenue from re-directed supply}} - \underbrace{\gamma\cdot(1-\alpha)[Q_i - y(p_i,p_j,\lambda)-\epsilon_i])}_{\text{underage cost}}$$

if $\epsilon_i < Q_i - y(p_i,p_j,\lambda)$. Here, $Q_i - y(p_i,p_j,\lambda) - \epsilon_i$ denotes the amount of shortage faced by buyer i, which is equal to the amount of rejected supply at buyer j given $\epsilon_i = -\epsilon_j$ and $Q_i + Q_j = 2$.

Let $\bar{\epsilon}_i(p_i,p_j,\lambda)$ denote $\mathbb{E}[\epsilon|\epsilon < Q_i - y(p_i,p_j,\lambda)]$ and using the shorthand $y_i := (p_i,p_j,\lambda)$, $\bar{\epsilon}_i := \bar{\epsilon}_i(p_i,p_j,\lambda)$, buyer i's expected profit can be written as $\pi_i(p_i,p_j,\lambda) = [1-F(Q_i - y_i;\sigma)]\cdot Q_i(1-p_i) + F(Q_i - y_i;\sigma)\cdot [(y_i+\bar{\epsilon}_i)(1-p_i) + \alpha[Q_i - y_i - \bar{\epsilon}_i](1-p_i) - \gamma\cdot(1-\alpha)[Q_i - y_i - \bar{\epsilon}_i])]$. This allows us to compute the competitive equilibrium by numerically solving for the fixed point.

Figure 12 illustrates market dynamics for various values of α under the competitive equilibrium. Overall, welfare benefits to both sellers (Figure 12a) and buyers (Figure 12c) are still present, and the amount of wasted supply that goes to outside options decreases with increasing transparency (Figure 12b). When α is large, however, price competition in the market significantly decreases, and changes in buyer profit becomes smaller. Intuitively, since both buyers are able to recover a sizable amount of lost supply from each other, there is less need to compete for supply or worry about any shortage. This also reduces profit gains from demand signaling.

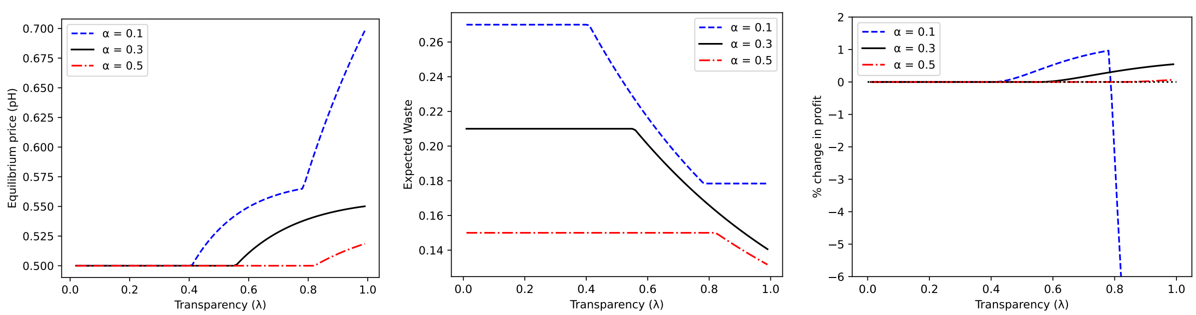

(a) p_H^* v.s. λ (b) Expected waste v.s. λ (c) % Chg in $\bar{\pi}(\mathbf{p}^*,\lambda)$ v.s. λ

Figure 12 **Buyers turn down additional supply and a fraction α of excess supply is sold to the other buyer. $\epsilon \sim U[-\sigma,\sigma]$ and $\epsilon_i = -\epsilon_j$. The other parameters are $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.25, \Delta = 0.3, Q_H + Q_L = 2$.**

A.3. Explicit Models of Supply Uncertainty

In this section, we provide several explicit models of supply uncertainty, motivated by the two main sources of uncertainty in the market. First, individual supply quantities from sellers vary due to yield uncertainty and harvest decisions. Second, sellers may consider factors other than prices when deciding whom to trade with. Figure 13 shows fluctuations in weekly supply from a seller on the *PemPem* platform and demonstrates that both uncertainty sources are present. The amount supplied by the seller varies significantly every week, reflecting yield uncertainty. In addition, while the seller typically trades with one main buyer (green), he occasionally trades away with others. When he trades away, this can be due to a higher prices (orange) or other reasons (purple).

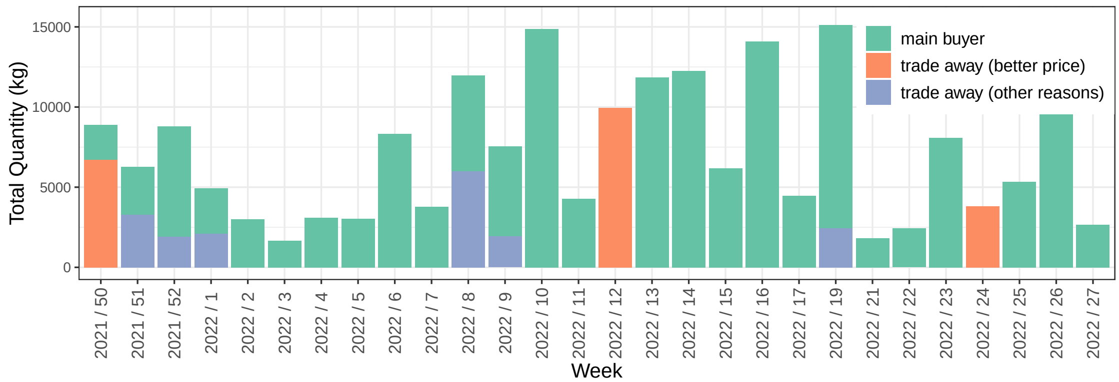

Figure 13 **Weekly supply quantities of one seller over time. Green represents supply sold to the main buyer, orange and purple represent supply sold to others that can and cannot be explained by a higher price respectively.**

We now present models that capture these two sources of uncertainty. Under each model, we also calculate the equilibrium prices and profits by numerically solving for the fixed point in the competitive equilibrium (see Figures 14, 15, 16).

To see how individual seller decision and yield can lead to supply uncertainty faced by buyers, we discretize the market into M sellers, indexed by m. Individual supply by seller m is denoted as v_m.

A.3.1. Normal yield uncertainty. We consider first two models where each seller has uncertain yield. For tractability, we capture yield uncertainty using a normal distribution. For both Model 1(a) and 1(b), we normalize each seller's yield uncertainty so that the total market supply follows the distribution $\mathcal{N}(2, \sigma^2)$.

- **Model 1(a) IID yield uncertainty**. Each period, individual seller's yield follows the distribution

$$v_m \overset{iid}{\sim} \mathcal{N}\left(\frac{2}{M}, \frac{\sigma^2}{M}\right).$$

  Yield uncertainty is independently distributed among sellers, and so total market supply follows the distribution $\sum_{m=1}^{M} v_m \sim \mathcal{N}(2, \sigma^2)$. Each seller decides on who to sell to as described in Section 2.2. Under our base model with a continuum of sellers, a fraction of $\frac{y(p_i,p_j,\lambda)}{2}$ of the sellers will sell to buyer i, leading to an expected supply of $y(p_i, p_j, \lambda)$. Analogously, in the discretized market, for buyer i, the number of sellers selling to him is approximately $n = \text{round}\left(y(p_i, p_j, \lambda) \cdot \frac{M}{2}\right)$. Thus, total supply procured by buyer i follows the distribution of the sum of n iid normal random variables, which yields $\sum_{m=1}^{n} v_m \sim \mathcal{N}\left(\frac{2n}{M}, \frac{n}{M}\sigma^2\right)$. Using continuous approximation $n \approx y(p_i, p_j, \lambda) \cdot \frac{M}{2}$, buyers face supply uncertainty given by

$$y(p_i, p_j, \lambda) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}\left(0, \frac{y(p_i, p_j, \lambda)}{2}\sigma^2\right),$$

where ϵ_i, ϵ_j are *independent*.

The scale parameter is given by $\sigma(y) = \sigma_0 \cdot \sqrt{y}$, where $\sigma_0 := \sigma/\sqrt{2}$. It is increasing but not linear in y. Alternatively, this form of uncertainty can also be rewritten in multiplicative form as $y(p_i, p_j, \lambda) + \sqrt{y(p_i, p_j, \lambda)}\varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma^2/2)$ has a constant scale parameter, $\sigma/\sqrt{2}$.

- **Model 1(b) Perfectly correlated yield uncertainty**. Each period, each seller has perfectly correlated yield quantity that follows the distribution $v_1 = v_2 = \cdots = v_m \sim \mathcal{N}\left(\frac{2}{M}, \frac{\sigma^2}{M^2}\right)$. Again, this yields total market supply following the distribution $\sum_{m=1}^{M} v_m \sim \mathcal{N}(2, \sigma^2)$. Similar to Model 1(a), we use the continuous approximation for the number of sellers that sell to buyer i: $n \approx y(p_i, p_j, \lambda) \cdot \frac{M}{2}$. Consequently, buyers face perfectly correlated supply uncertainty given by

$$y(p_i, p_j, \lambda) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}\left(0, \frac{y(p_i, p_j, \lambda)^2}{4}\sigma^2\right),$$

where ϵ_i, ϵ_j are *perfectly correlated*, $\epsilon_i = \epsilon_j$.

The scale parameter is given by $\sigma(y) = \sigma_0 \cdot y$, where $\sigma_0 := \sigma/2$. It is linear in y. Alternatively, this form of uncertainty can also be rewritten in multiplicative form as $y(p_i, p_j, \lambda) + y(p_i, p_j, \lambda)\varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma^2/4)$ has a constant scale parameter, $\sigma/2$.

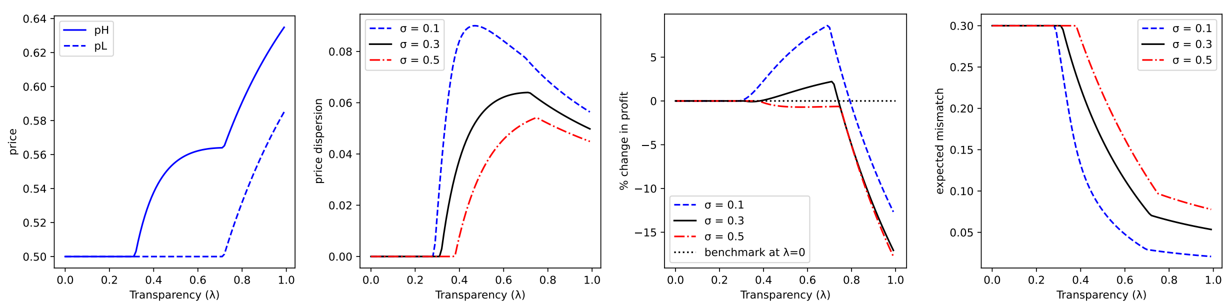

(a) $p^*(\lambda)(\sigma = 0.3)$ (b) $p^*_\Delta(\lambda)$ v.s. λ (c) profit (d) mismatch

Figure 14 **Model (1a) with parameters** $\Delta = 0.3, \underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2$.

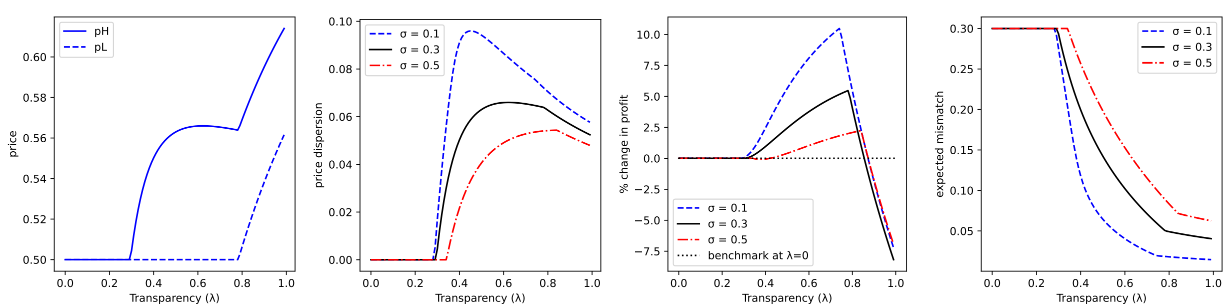

(a) $p^*(\lambda)(\sigma = 0.3)$ (b) $p^*_\Delta(\lambda)$ v.s. λ (c) profit (d) mismatch

Figure 15 **Model (1b) with parameters** $\Delta = 0.3, \underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2$.

A.3.2. Random selling decisions. Consider a market with M sellers. Uninformed sellers randomly select one of the buyers to sell to, while informed sellers always sell to the buyer providing higher utility. In other words, we assume that uninformed sellers make random decisions, while the decisions from informed sellers are fully rational.

The total number of uninformed sellers in the market is approximately $n^U = \text{round}(M(1-\lambda))$, and the number of informed sellers selling to buyer i is approximately $n^I = \text{round}(\frac{M}{2} \cdot [y(p_i, p_j, \lambda) - (1-\lambda)])$.

- **Model (2a) Random decisions without yield uncertainty.** Suppose each seller offers a fixed supply of $\frac{2}{M}$ each period, such that total market supply is constantly 2. Then, total supply procured by buyer i each period follows the distribution $\frac{2}{M} \cdot [\text{Binomial}(n^U, 1/2) + n^I)]$.

  We use normal approximation for binomial distributions and continuous approximation for n^I, n^U to simplify buyer's supply distribution. This yields additive supply uncertainty, given by

  $$y(p_i, p_j, \lambda) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}\left(0, (1-\lambda)/M\right),$$

  where ϵ_i, ϵ_j have *perfect negative correlation*, $\epsilon_i + \epsilon_j = 0$.

  The scale parameter is given by $\sigma(\lambda) = \sqrt{\frac{1-\lambda}{M}}$, which is decreasing in transparency λ and independent on expected supply $y(p_i, p_j, \lambda)$.

- **Model (2b) Random decisions with yield uncertainty.** Suppose that, each period, each seller also has i.i.d. supply quantity $v_m \sim \mathcal{N}\left(2/M, \sigma^2/M\right)$, as in Model 1(a). Then, each buyer receives supply that is the sum of a random binomial number of normal random variables, $\sum_{m=1}^{\tilde{n}^U + n^I} v_m$, where $\tilde{n}^U \sim \text{Binomial}(n^U, 1/2)$. While the exact form of supply distribution is complicated, it is easy to verify that supply uncertainty faced by individual buyers can be approximately written as $y_i + \epsilon_i$, where ϵ_i has 0 mean with variance both decreasing in transparency λ and increasing expected supply, $y(p_i, p_j, \lambda)$.

  An alternative and more tractable way to model yield uncertainty under this setting is to consider sellers' supply as binary: each period, each seller is actively selling w.p. ρ. Binary supply could capture yield uncertainty due to random harvest schedules of individual farmers. For instance, in the palm fruit market, each farmer typically harvests fruit for sale once every week or once every two weeks.

  When active, each seller has a constant supply of $\frac{2}{\rho M}$, i.e., $v_m \overset{iid}{\sim} \frac{2}{\rho M}\text{Bernoulli}(\rho)$ and $\mathbb{E}[v_m] = \frac{2}{M}$. Then, total supply in the market follows the distribution $\sum_{m=1}^{M} v_m \sim \frac{2}{\rho M} \cdot \text{Binomial}(M, \rho)$. Note that we have calibrated the distribution of v_m such that total market supply has a normalized expectation of 2. Consequently, smaller ρ corresponds to higher yield uncertainty, and vice versa. Model 2(a) is a special case of 2(b) with $\rho = 1$.

  Supply procured by buyer i now follows the distribution $\frac{2}{\rho M} \cdot [\text{Binomial}(n^U, \rho/2) + \text{Binomial}(n^I, \rho)]$. Using normal approximation for binomial distributions and continuous approximation for n^I, n^U, buyers face supply uncertainty with marginal distribution,

  $$y_i + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}\left(0, \frac{4}{\rho M}\left[(1-\lambda)(1-\rho/2)/2 + \frac{1}{2}[y(p_i, p_j, \lambda) - (1-\lambda)](1-\rho)\right]\right),$$

  where ϵ_i, ϵ_j are *negatively correlated.* Moreover, the uncertainty term can be equivalently separated as

  $$y_i + \epsilon_i^{(\lambda)} + \epsilon_i^{(y)}, \quad \epsilon_i^{(\lambda)} \sim \mathcal{N}(0, (1-\lambda)/M) \quad \epsilon_i^{(y)} \sim \mathcal{N}\left(0, y(p_i, p_j, \lambda) \cdot \frac{2(1-\rho)}{M\rho}\right).$$

Observe that $\epsilon_i^{(\lambda)}$ captures uncertainty from random decisions by uninformed sellers (same as that in Model 2(a)), which decreases with λ. Meanwhile, $\epsilon_i^{(y)}$ captures independent yield uncertainty from individual sellers, with a scale parameter proportional to $\sqrt{y(p_i, p_j, \lambda)}$, similar to that in Model 1(a).

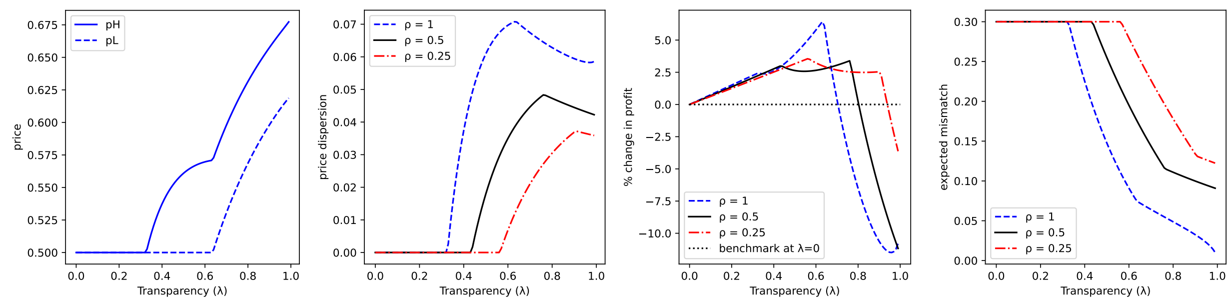

(a) $p^*(\lambda)(\rho = 1)$ (b) $p^*_\Delta(\lambda)$ v.s. λ (c) profit (d) mismatch

Figure 16 **Model 2 with parameters $\Delta = 0.3, \underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2$. The blue-dotted line with $\rho = 1$ represents Model 2(a) with no yield uncertainty. The other two lines represent Model 2(b) for $\rho = 0.5$ (lower yield uncertainty) and $\rho = 0.25$ (higher yield uncertainty).**

Lastly, Figure 17 confirms graphically that buyer profit remains concave under yield-uncertainty models, Model 1(a), 1(b) and 2(b), thereby satisfying Assumption 3.

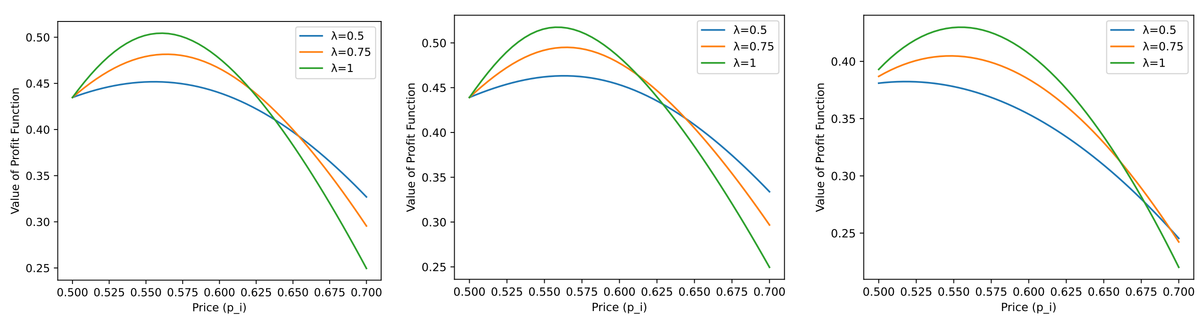

(a) Model 1(a) (b) Model 1(b) (c) Model 2(b)

Figure 17 **The profit function, $\pi_i(p_i, p_j, \lambda)$ v.s. p_i. It remains concave under yield-dependent uncertainty models, Model 1(a), 1(b) and 2.**

A.4. Discussion of the Uniform Reservation Price Assumption

This section presents additional results and numerical simulations in markets with non-uniform reservation prices. We show that replacing the uniform reservation price assumption has very limited impact on the qualitative insights presented in the paper.

A natural way that variations in reservation prices can arise is that each seller has lower reservation price at the buyer closer to him compared to the buyer further away. This can be modeled via reservation *utility* of sellers, taking into account both the selling price and transportation costs. Let r be the uniform reservation utility for all sellers. That is, they desire to sell all of their produce if the per-unit utility is at or above r, and will sell none of it below r. Since sellers' per-unit utility is defined by price minus transportation cost, a seller located at x with reservation utility r will equivalently have a reservation price of $\underline{p}(x) = r + tx$ at one buyer and $\underline{p}(x) = r + t(1-x)$ at the other.

We assume that 'loyalty' is purely based on transportation costs (i.e. sellers on $[0, 1/2]$ sell to buyer 0, and on $[1/2, 1]$ sell to buyer 1). Then, the expected supply $y(p_i, p_j, \lambda)$ given in Equation (1) becomes

$$y'(p_i, p_j, \lambda) = \begin{cases} 0 & \text{if } p_i < r, \\ 2(1-\lambda) \cdot \min\left\{\frac{1}{2}, \frac{p_i - r}{t}\right\} + 2\lambda \cdot \min\left\{\left(\frac{1}{2} + \frac{p_i - p_j}{2t}\right)^+, \frac{p_i - r}{t}\right\} & \text{if } r \leqslant p_i < r + t, \\ (1-\lambda) + 2\lambda \cdot \min\left\{\left(\frac{1}{2} + \frac{p_i - p_j}{2t}\right)^+, 1\right\} & \text{if } p_i \geqslant r + t. \end{cases}$$

Using $y'(p_i, p_j, \lambda)$ above, we numerically solve for $\mathbf{p}^*(\lambda)$ and $\mathbf{p}^{c*}$ in a market with uniform reservation utility, using the same set of market parameters used in all illustrations in the paper. We use $r = 0.4, t = 0.2$ versus a uniform reservation price of $\underline{p} = 0.5$ for numerical calculations.

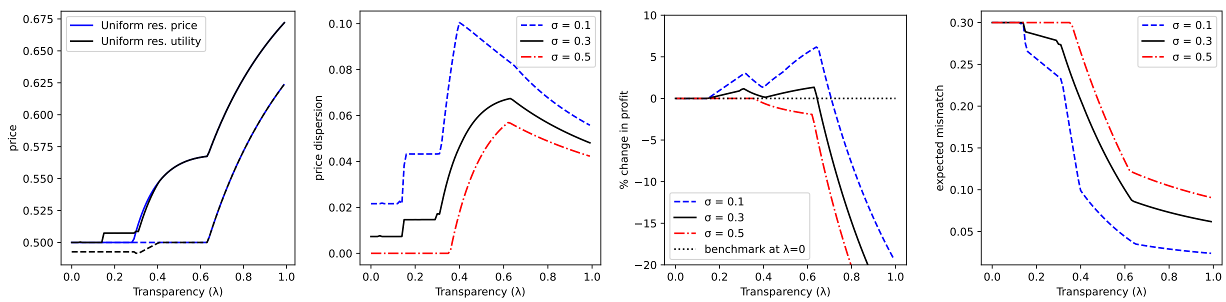

(a) $p^*(\lambda)(\sigma = 0.3)$ (b) $p^*_\Delta(\lambda)$ v.s. λ (c) profit (d) mismatch $(-s_H)$

Figure 18 **Competitive market with uniform reservation utility.** $\Delta = 0.3, \underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2, r = 0.4.$

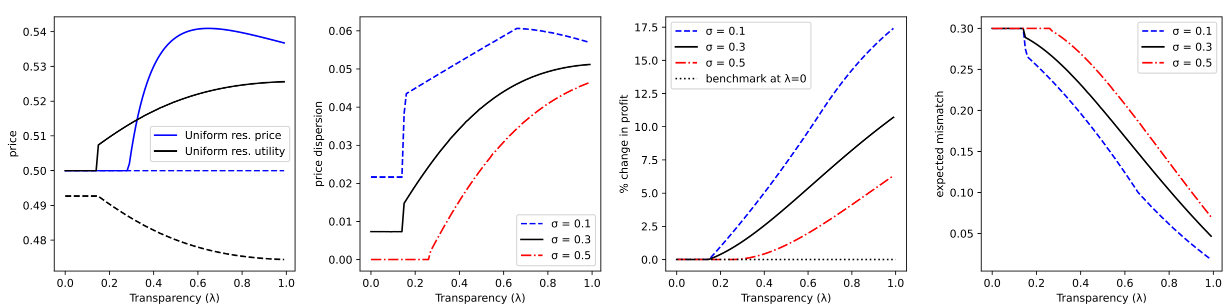

(a) $p^*(\lambda)(\sigma = 0.3)$ (b) $p^*_\Delta(\lambda)$ v.s. λ (c) profit (d) mismatch $(-s_H)$

Figure 19 **Collusive market with general reservation prices.** $\Delta = 0.3, \underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.2, r = 0.4.$

Competitive Equilibrium. Figure 18 illustrates the competitive equilibrium $\mathbf{p}^*(\lambda)$ with uniform reservation utility. We observe that market dynamics (e.g. prices, existence of Pareto improvement and improved efficiency) closely resemble those observed under uniform reservation price (see Figure 4).

Collusive Equilibrium. Figure 19 illustrates the profit-maximizing collusive equilibrium $\mathbf{p}^{c*}$ with uniform reservation utility. We observe that while most market dynamics remain similar, the low-demand buyer's may lower his price with increased transparency under uniform reservation utility, instead of keeping it constant. This raises the question whether sellers' welfare are always protected under buyer collusion. The following proposition shows that this is indeed the case, and increased transparency always guarantee at least *weak* Pareto improvement in the market. The result requires that sellers are loyal to the closest buyer.

Proposition 8. *Suppose sellers are uniformly distributed on* $[0,1]$ *and loyal to the closest buyer. Then, if sellers have uniform reservation utility, transparency does not reduce average market prices under jointly profit-maximizing collusion and always leads to at least weak Pareto improvement in the market.*

The intuition behind this result is that, with increased transparency, it becomes worthwhile for buyers with high demand to increase their prices to attract informed sellers that are located further away. Similarly, it may be optimal to lower the price under low demand. As demonstrated in Figure 20, this price-adjusting effect is generally either symmetric or skewed towards a price increase when transparency is sufficiently high, or the underage costs γ are high. In particular, Figure 20b shows an example where average market prices show significant increase for some λ, leading to strict improvements in utility for all sellers. The proof of Proposition 8 can be found in Appendix E.

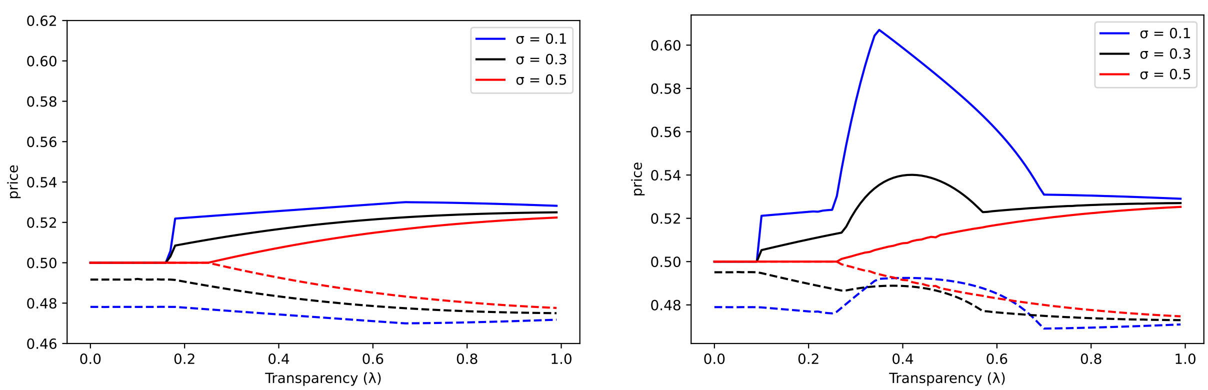

(a) $\gamma = 0.1, h = 0.5$ (b) $\gamma = 0.6, h = 0.5$

Figure 20 $\mathbf{p}^{c*}$ **with uniform reservation utility. Dotted lines indicate** p_L^{c*}**, and solid lines indicate** p_H^{c*}**.**

A.5. General Forms of Asymmetric Demand

So far, we have only considered scenarios where buyers are in the opposite demand states every period, and $Q_H + Q_L = 2$ so that total supply and total demand are balanced in expectation. In this section, we relax both assumptions and consider the general case where the demand states are arbitrarily correlated, and allow that $Q_H + Q_L \neq 2$.

Under this general setting, both buyers aim to fulfill regular, recurring procurement targets each period. Different demand states arise from delivery timelines, which can be asynchronized among the two buyers (see Section 2.3). Each buyer is in high demand Q_H for a $\eta > 0$ fraction of periods (for example, one day every week right before the delivery deadline), which is equal among the two buyers. Each buyer is in low demand state Q_L for the remaining $1-\eta$ fraction of periods. The demand states among the two buyers can be arbitrarily correlated, and we pose no restriction on possible values of Q_L, Q_H besides that $Q_L < Q_H$ such that $F(1-Q_L;\sigma) > F(1-Q_H;\sigma)$.

We let $\tau \leqslant \eta$ denote the fraction of periods that both buyers are simultaneously in high demand. τ varies with delivery timelines and demand correlation as follows.

- When buyers are never simultaneously in high demand (e.g. when the buyers' delivery deadlines are completely unaligned), we have $\tau = 0$ and $\eta \leqslant 0.5$. In this case, during any period, buyers are in opposite demand states with probability 2η, and are simultaneously in low demand with probability $1-2\eta$.
- Alternatively, demand may be positively correlated such that there are some periods where buyers are simultaneously in high demand, $\tau > 0$. Then, during any period, buyers are in opposite demand states with probability $2(\eta-\tau)$, are simultaneously in low demand with probability $1-2\eta+\tau$, and are simultaneously in high demand with probability τ.
- When demands are perfectly correlated (e.g. when both buyers have the same delivery deadlines), $\tau = \eta$, and buyers are never in opposite demand states. During any period, buyers are simultaneously in low demand with probability $1-\tau$, and are simultaneously in high demand with probability τ.

Analogous to Equation (8), buyer welfare $\bar{\pi}$ is defined as the expected profit averaged over demand states,

$$\bar{\pi} = \tau \cdot \pi_{HH} + (\eta - \tau)\cdot(\pi_{HL} + \pi_{LH}) + (1-2\eta+\tau)\cdot \pi_{LL},$$

and strong Pareto improvement is defined in Equation (10). We used the suffix HH, LL, HL, LH to indicate the four different scenarios of demand states in the market respectively.

First, recall from Lemma 2 that individual buyer can only strictly benefit if the expected supply-demand mismatch is reduced. When buyers are in the same demand state (either HH or LL), however, buyers will set equal prices in the competitive equilibrium, and thus supply-demand mismatch remains unchanged compared to the initial market where both buyers price at $\underline{p}$. Hence, buyers never strictly benefit from transparency when both are simultaneously in high or low demand. Achieving two-sided benefits thus requires that $\tau < \eta$, such that buyers spend sometime in the opposite demand states.

Theorem 5 establishes that the results in Theorem 2 hold in this general setting as long as buyers are not often simultaneously in the high demand state. Note that it assumes the buyers observe each others' demand states at the time of pricing.

Theorem 5. *Suppose buyers observe each others' demand states, $Q_L < Q_H$ such that $F(1-Q_L;\sigma) > F(1-Q_H;\sigma)$, and buyers are simultaneously in demand state Q_H for a τ fraction of periods. Consider any market satisfying the competitive market condition, i.e., $\hat{\lambda} < 1$ and $\lambda > \hat{\lambda}$, where $\hat{\lambda} = \frac{t}{1-\underline{p}+\gamma-(\gamma+h)F(1-Q_H;\sigma)}$.*

Then, if $1-Q_L > \underline{\Delta}$ and $h \geqslant 1-\underline{p}$, there always exists $0 < \bar{\tau} < \eta$ such that for all $\tau \in [0, \bar{\tau})$, there exists a threshold $\lambda^*(\tau)$ such that any $\lambda \in (\hat{\lambda}, \lambda^*(\tau))$ is strongly Pareto improving under the competitive equilibrium. Herein, $\underline{\Delta}$ is the smallest value such that $F(\underline{\Delta};\sigma) = \frac{\gamma+1-\underline{p}}{\gamma+h}$.

The conditions on λ, Δ and h required under Theorem 5 are similar to those identified by Theorem 2. In addition, Theorem 5 demands that the buyers are not often simultaneously in high demand state, i.e., $\tau < \bar{\tau}$. Intuitively, when both are in high demand, price transparency brings no benefit of demand signaling, and intense price competition further harms buyers' profit. In practice, this suggests that two-sided benefits are likely in markets where procurement delivery deadlines occur at different times, or quantity targets are negatively correlated among the buyers. If the contractual timelines or quantities are highly correlated, however, the benefits of price transparency for buyers will be limited.

Figures 21 and 22 illustrate market dynamics under overall high demand, $Q_L + Q_H > 2$ versus overall low demand, $Q_L + Q_H < 2$. The equilibrium prices when buyers are in opposite demand states are computed by numerically solving for the fixed-point of the equilibrium. When overall demand is greater than market supply, intense market competition means that even when two-sided benefits arise, there is very limited range of transparency beneficial to the buyers (see the blue dotted line in Figure 21b). Conversely, when overall demand is lower than market supply, lower market competitive means greater buyer welfare (see Figure 22b).

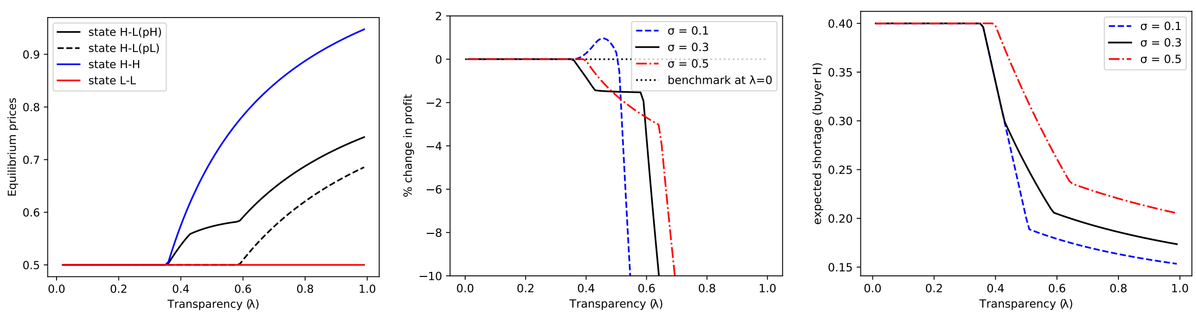

(a) $p^*(\lambda)$ v.s. λ, $\sigma = 0.3$. (b) % Chg in $\bar{\pi}(\mathbf{p}^*, \lambda)$ v.s. λ (c) $Q_H - y_H$ v.s. λ

Figure 21 **Overall high demand (supply shortage):** $Q_H = 1.4, Q_L = 0.8$. $\epsilon \sim U[-\sigma, \sigma]$. **The other parameters are** $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.25$.

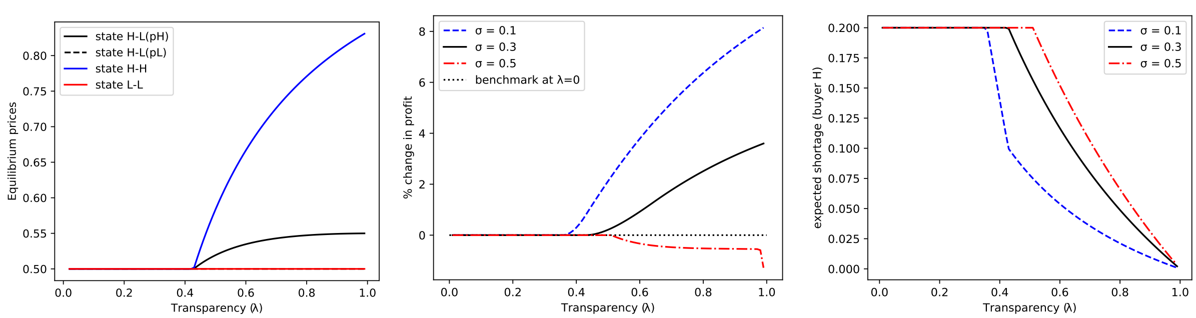

(a) $p^*(\lambda)$ v.s. λ, $\sigma = 0.3$. (b) % Chg in $\bar{\pi}(\mathbf{p}^*, \lambda)$ v.s. λ (c) $Q_H - y_H$ v.s. λ

Figure 22 **Overall low demand (supply surplus):** $Q_H = 1.2, Q_L = 0.6$. $\epsilon \sim U[-\sigma, \sigma]$. **The other parameters are** $\underline{p} = 0.5, \gamma = 0.2, h = 0.5, t = 0.25$.

A.6. Impact of Price Discrimination under Partial Transparency

When $\lambda < 1$, maintaining partial transparency requires that price information cannot be shared between smallholder farmers. Buyers could then set a competitive price on the platform but charge the reservation price $\underline{p}$ for uninformed sellers offline. Our modeling framework can be extended to this scenario. Although full analysis of market welfare under price discrimination is out of the scope of this paper, we provide a high-level discussion of potential impact of price discrimination in this section, while leaving a more thorough analysis to future work.

Consider buyers that engage in price discrimination such that they always charge uninformed seller the reservation price, $\underline{p}$. Buyer's new riskless profit becomes $(1-\lambda)(1-\underline{p}) + [y(p_i, p_j, \lambda) - (1-\lambda)](1-p_i)$, where the first term captures expected profit from uninformed sellers, and the second term captures expected profit from informed sellers. The remainder of the profit function remains unchanged.

The buyers then decide on their pricing for informed sellers. It is easy to verify via the first order conditions that when $p_H^*, p_L^* > \underline{p}$, competitive equilibrium prices must satisfy

$$\frac{1}{2}(p_H^* + p_L^*) = 1 + \frac{\gamma - h}{2} - t, \quad p_H^* - p_L^* = \frac{\gamma + h}{3}\left[1 - 2F\left(\frac{\lambda}{t}p_\Delta^*(\lambda) - \Delta; \sigma\right)\right]. \tag{15}$$

for informed sellers, and $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ for uninformed sellers.

We compare the market prices in Equation (15) with those in Equations (6) and (7) under Theorem 1, and make the following observations.

From the perspective of seller welfare, sellers without access to the platform never benefit from improved prices. Furthermore, market prices for informed sellers are higher with price discrimination than without it. This suggests that price discrimination can worsen inequality between informed and uninformed sellers. Therefore, two-sided benefits occur only when all sellers have equal access to the platform (which may then offer partial transparency to everyone).

From the perspective of buyers welfare, we observe a similar effect of demand signaling through price differentiation. Price dispersion characterized in Equation (15) remains identical to (7), and is decreasing in transparency λ and increasing in supply uncertainty, σ. At the same time, price discrimination can lead to more intense price competition for supply from informed sellers. As a result, it is not immediately clear whether and when buyers will be motivated to engage in price discrimination for greater profits.

Additional numerical simulations further suggest that with careful selection of transparency levels and equal access for all sellers, two-sided benefits can still be achieved in the presence of price discrimination. Identifying the best approaches to discourage price discrimination and ensure two-sided benefits presents a meaningful question for future work.

Appendix B: Notation and Expected Buyer Profit

Throughout the appendix we will use the following notation.

Market prices: Market prices are defined either by p_H, p_L or by $\bar{p}, p_\Delta$, where $\bar{p} = \frac{1}{2}(p_H + p_L)$ is the average price, and $p_\Delta = p_H - p_L$ is the price dispersion. When there is no ambiguity in context, we will use the short form $\mathbf{p}$ to represent market prices.

Safety stock: The safety stock of buyer i, denoted as $s_i(p_i, p_j, \lambda)$, is defined as the difference between the expected supply and contractual demand,

$$s_i(p_i, p_j, \lambda) = y(p_i, p_j, \lambda) - Q_i \quad \text{for} \quad i \in \{H, L\}.$$

Note that $s_i(p_i, p_j, \lambda)$ depends on market prices only via p_Δ, and when the dependence on p_Δ is evident we omit it for concision:

$$s_H(\lambda) = -s_L(\lambda) = \frac{\lambda}{t}\max\{-t, \min\{p_\Delta, t\}\} - \Delta.$$

In Section 4.1, we also introduced the notion of market efficiency measured by the absolute level of safety stock, or equivalently mismatch between demand and expected supply: $s(\mathbf{p}, \lambda) := |s_H(\lambda)| = |s_L(\lambda)|$.

B.1. Expected Buyer Profit

In this section, we derive closed-form expressions for buyer expected profit. Expressions for profit labelled '(P-X)' (where different values of X correspond to different cases in the analysis) will be used in our proofs.

B.1.1. Case 1: $|p_i - p_j| \leqslant t$. When $|p_i - p_j| \leqslant t$, $s_H(\lambda) = -s_L(\lambda) = \frac{\lambda}{t}p_\Delta - \Delta$. The expected profit of each buyer given the market prices and λ is

$$\pi_i(p_i, p_j, \lambda) \equiv \pi_i(s_i(\lambda), p_j, \lambda) = \Psi(p_i, p_j, \lambda) - L_i(p_i, p_j, \lambda), \tag{P-1}$$

where $s_i(\lambda) = \frac{\lambda}{t}(p_i - p_j) - \Delta$ is the safety stock,

$$\Psi(p_i, p_j, \lambda) = y(p_i, p_j, \lambda)(1 - p_i) = (Q_i + s_i(\lambda))\left[1 - p_j - (s_i(\lambda) + Q_i - 1)\cdot\frac{t}{\lambda}\right]$$

is the riskless profit, and

$$L(p_i, p_j, \lambda) = \gamma\mathbb{E}\left[(s_i(\lambda) + \epsilon)^-\right] + h\mathbb{E}\left[(s_i(\lambda) + \epsilon)^+\right] = \gamma\int_{-\infty}^{-s_i(\lambda)} F(\epsilon;\sigma)d\epsilon + h\int_{-s_i(\lambda)}^{\infty} 1 - F(\epsilon;\sigma)d\epsilon$$

is the expected loss from uncertain supply due to underage and overage costs. Observe that the expected overage can be equivalently written as $h\int_{-s_i(\lambda)}^{\infty} 1 - F(\epsilon;\sigma)d\epsilon = h\int_{-\infty}^{s_i(\lambda)} F(\epsilon;\sigma)d\epsilon = h\left[\int_{-\infty}^{-s_i(\lambda)} F(\epsilon;\sigma)d\epsilon + s_i(\lambda)\right]$. The first equality follows from a change of variable from ϵ to $-\epsilon$ and that $1 - F(\epsilon;\sigma) = F(-\epsilon;\sigma)$ by symmetry. The second equality uses the fact that $\int_{-\infty}^{s_i(\lambda)} F(\epsilon;\sigma)d\epsilon = \int_{-\infty}^{-s_i(\lambda)} F(\epsilon;\sigma)d\epsilon + \int_{-s_i(\lambda)}^{s_i(\lambda)} F(\epsilon;\sigma)d\epsilon$ and $\int_{-s_i(\lambda)}^{s_i(\lambda)} F(\epsilon;\sigma)d\epsilon = s_i(\lambda)$ by symmetry.

The expression for each buyer's expected profit implies that the first order condition (FOC) of the best response for each buyer is

$$\frac{\partial\pi_i(s_i(\lambda), p_j, \lambda)}{\partial s_i(\lambda)} = -2s_i(\lambda)\cdot\frac{t}{\lambda} + 1 - p_j + \frac{t}{\lambda}(1 - 2Q_i) + \gamma[1 - F(s_i(\lambda);\sigma)] - hF(s_i(\lambda);\sigma) = 0. \tag{P-FOC}$$

The expected individual and joint profits for any given λ can also be written in terms of the average price $\bar{p}$ and price dispersion p_Δ as follows,

$$\pi_H(\bar{p}, p_\Delta, \lambda) = (1 + \tfrac{\lambda}{t} p_\Delta)(1 - \bar{p} - p_\Delta/2) - \gamma \textstyle\int_{-\infty}^{\Delta - \frac{\lambda}{t} p_\Delta} F(\epsilon; \sigma) d\epsilon - h \left[\textstyle\int_{-\infty}^{\Delta - \frac{\lambda}{t} p_\Delta} F(\epsilon; \sigma) d\epsilon - \Delta + \tfrac{\lambda}{t} p_\Delta \right], \quad \text{(P-H)}$$

$$\pi_L(\bar{p}, p_\Delta, \lambda) = (1 - \tfrac{\lambda}{t} p_\Delta)(1 - \bar{p} + p_\Delta/2) - \gamma \textstyle\int_{-\infty}^{-\Delta + \frac{\lambda}{t} p_\Delta} F(\epsilon; \sigma) d\epsilon - h \left[\textstyle\int_{-\infty}^{-\Delta + \frac{\lambda}{t} p_\Delta} F(\epsilon; \sigma) d\epsilon + \Delta - \tfrac{\lambda}{t} p_\Delta \right], \quad \text{(P-L)}$$

$$\begin{aligned} \pi_H(\bar{p}, p_\Delta, \lambda) + \pi_L(\bar{p}, p_\Delta, \lambda) &= 2 - 2\bar{p} - \frac{\lambda}{t} p_\Delta^2 - (\gamma + h) \left[2 \int_{-\infty}^{\Delta - \frac{\lambda}{t} p_\Delta} F(\epsilon; \sigma) d\epsilon - \left(\Delta - \frac{\lambda}{t} p_\Delta \right) \right] \\ &\equiv 2 - 2\bar{p} - \frac{\lambda}{t} p_\Delta^2 - (\gamma + h) \left[2 \int_{-\infty}^{-(\Delta - \frac{\lambda}{t} p_\Delta)} F(\epsilon; \sigma) d\epsilon + \Delta - \frac{\lambda}{t} p_\Delta \right] \end{aligned} \quad \text{(P-joint)}$$

The last expression uses $\int_{-\infty}^{-(\Delta - \frac{\lambda}{t} p_\Delta)} F(\epsilon; \sigma) d\epsilon = \int_{-\infty}^{\Delta - \frac{\lambda}{t} p_\Delta} F(\epsilon; \sigma) d\epsilon - \Delta + \frac{\lambda}{t} p_\Delta$ from symmetry of $F(\epsilon; \sigma)$.

Similarly, each buyer's expected profit $\bar{\pi}(\mathbf{p}, \lambda)$ can be written as $\bar{\pi}(\mathbf{p}, \lambda) = \bar{\Psi}(\mathbf{p}, \lambda) - \bar{L}(s(\mathbf{p}, \lambda), \lambda)$, where the riskless profit and expected loss are

$$\bar{\Psi}(\mathbf{p}, \lambda) = \frac{1}{2} \sum_{i \in \{H, L\}} y(p_i, p_j, \lambda)(1 - p_i) = 1 - \bar{p} - \frac{\lambda}{2t} p_\Delta^2, \quad \text{(RP)}$$

$$\bar{L}(s(\mathbf{p}, \lambda), \lambda) = (\gamma + h) \left[\int_{-\infty}^{-s(\mathbf{p}, \lambda)} F(\epsilon; \sigma(\lambda)) d\epsilon + \frac{s(\mathbf{p}, \lambda)}{2} \right]. \quad \text{(L)}$$

B.1.2. Case 2: $|p_i - p_j| \geqslant t$. When $|p_i - p_j| \geqslant t$, the expected supply $y(p_i, p_j, \lambda)$ becomes

$$y(p_i, p_j, \lambda) = \begin{cases} 1 + \lambda, & \text{if } p_i - p_j \geqslant t, \\ 1 - \lambda, & \text{if } p_i - p_j \leqslant -t, \end{cases}$$

and the safety stock is $s_i(\lambda) = 1 \pm \lambda - Q_i$. The expected profits are, denoting $L_i(p_i, p_j, \lambda) \equiv L_i(s_i, \lambda)$,

$$\pi_i(p_i, p_j, \lambda) = \begin{cases} (1 + \lambda)(1 - p_i) - L_i(1 + \lambda - Q_i, \lambda) & \text{if } p_i - p_j \geqslant t, \\ (1 - \lambda)(1 - p_i) - L_i(1 - \lambda - Q_i, \lambda) & \text{if } p_i - p_j \leqslant -t. \end{cases} \quad \text{(P-2)}$$

Note that the expected loss L is independent of price, and $\pi_i(p_i, p_j, \lambda)$ is decreasing in p_i in both cases. Therefore, $\pi_i(p_i, p_j, \lambda)$ is maximized at $p_i = p_j + t$ for all $p_i \geqslant p_j + t$. When $\lambda < 1$, $\pi_i(p_i, p_j, \lambda)$ is maximized at $p_i = \underline{p}$ for all $p_i \leqslant p_j - t$. In the special case where $\lambda = 1$, the profit $\pi_i(p_i, p_j, \lambda)$ is invariant for all $p_i \leqslant p_j - t$.

Appendix C: Full Characterization of the Competitive Equilibrium and Proof of Theorem 1

In this section, we provide a full characterization of equilibrium prices and prove Theorem 1.

C.1. Full Characterization of the Competitive Equilibrium

Let $p_i^{FOC}(p_j, \lambda)$ be defined by

$$p_i^{FOC}(p_j, \lambda) = p_j + \frac{t}{\lambda}(Q_i - 1) + \frac{t}{\lambda} \cdot s_i^{FOC}(\lambda), \quad \text{(FOC-1)}$$

where s_i^{FOC} satisfies

$$2 s_i^{FOC}(\lambda) \cdot \frac{t}{\lambda} + (\gamma + h) F(s_i^{FOC}(\lambda); \sigma) = 1 - p_j + \gamma + \frac{t}{\lambda}(1 - 2Q_i). \quad \text{(FOC-2)}$$

Let $\tilde{p}_H(\lambda), \tilde{p}_L(\lambda)$ denote the unique solution to the simultaneous equations

$$\frac{\tilde{p}_H(\lambda) + \tilde{p}_L(\lambda)}{2} = 1 + \frac{\gamma - h}{2} - \frac{t}{\lambda}, \quad \tilde{p}_H(\lambda) - \tilde{p}_L(\lambda) = \frac{\gamma + h}{3}(1 - 2F(s_H(\tilde{p}_H, \tilde{p}_L, \lambda); \sigma)). \quad \text{(FOC-3)}$$

Note that the equations in (FOC-3) are the same as Equations (6) and (7) in Theorem 1, and their unique solution $\tilde{p}_H(\lambda), \tilde{p}_L(\lambda)$ satisfies $p_H^{FOC}(\tilde{p}_L, \lambda) = \tilde{p}_H(\lambda)$, and $p_L^{FOC}(\tilde{p}_H, \lambda) = \tilde{p}_L(\lambda)$.

We state the unique characterization of equilibrium market prices $p_H^*(\lambda), p_L^*(\lambda)$ in the three mutually exclusive and collectively exhaustive cases below.

1. *Both buyers price at reservation price.* $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ if $(\gamma + h)F(\Delta;\sigma) - h \leqslant \frac{t}{\lambda} - (1 - \underline{p})$.
2. *Both buyers price above reservation price.* $p_H^*(\lambda) = \tilde{p}_H(\lambda), p_L^*(\lambda) = \tilde{p}_L(\lambda)$, and $p_H^*(\lambda) \geqslant p_L^*(\lambda) > \underline{p}$, if $\tilde{p}_L(\lambda) > \underline{p}$ and $(\gamma + h)F(\Delta;\sigma) - h > \frac{t}{\lambda} - (1 - \underline{p})$.
3. *One buyer prices at reservation price.* $p_H^*(\lambda) = p_H^{FOC}(\underline{p}, \lambda) > \underline{p}$, $p_L^*(\lambda) = \underline{p}$ if $\tilde{p}_L(\lambda) \leqslant \underline{p}$ and $(\gamma + h)F(\Delta;\sigma) - h > \frac{t}{\lambda} - (1 - \underline{p})$.

In the rest of this section, we prove that this gives a full characterization of equilibrium prices given that t exceeds some lower bound $\hat{t}$, and prove Theorem 1.

C.2. Best Response and First Order Conditions

We begin by characterizing the single-period best response of a buyer,

$$p_i^{BR}(p_j, \lambda) = \arg\max_{p_i \geqslant \underline{p}} \pi_i(p_i, p_j, \lambda), \quad \text{for} \quad i \in \{H, L\},$$

in terms of the first-order conditions (FOC-1), (FOC-2), and (FOC-3).

Lemma 3. *Define*

$$\tilde{p}_i^{BR}(p_j, \lambda) = \min\left\{\max\left\{\underline{p}, p_i^{FOC}(p_j, \lambda)\right\}, p_j + t\right\}$$

where $p_i^{FOC}(p_j, \lambda)$ satisfies the first-order conditions (FOC-1) *and* (FOC-2), *i.e.*

$$p_i^{FOC}(p_j, \lambda) = p_j + \frac{t}{\lambda}(Q_i - 1) + \frac{t}{\lambda} \cdot s_i^{FOC}(\lambda),$$

and s_i^{FOC} satisfies

$$2s_i^{FOC}(\lambda) \cdot \frac{t}{\lambda} + (\gamma + h)F(s_i^{FOC}(\lambda);\sigma) = 1 - p_j + \gamma + \frac{t}{\lambda}(1 - 2Q_i).$$

Then $\tilde{p}_i^{BR}(p_j, \lambda)$ is unique, and the best response p_i^{BR} to the other buyer's price p_j given any price transparency $\lambda \in (0, 1]$ is given by

(1) If $p_j - t \geqslant \tilde{p}_i^{BR}(p_j, \lambda) > \underline{p}$, $p_i^{BR}(p_j, \lambda) = \underline{p}$ for $\lambda < 1$ and can take any value in $[\underline{p}, p_j - t]$ at $\lambda = 1$.

(2) If $\tilde{p}_i^{BR}(p_j, \lambda) > p_j - t > \underline{p}$, $p_i^{BR}(p_j, \lambda) = \arg\max_{p_i \in \{\underline{p}, \tilde{p}_i^{BR}\}} \pi_i(p_i, p_j, \lambda)$.

(3) Otherwise, $p_i^{BR}(p_j, \lambda) = \tilde{p}_i^{BR}(p_j, \lambda)$.

In other words, buyer i's best response to price p_j is given by the first order conditions (FOC-1) and (FOC-2), restricted by boundary conditions which bound from below by $\underline{p}$ and from above by $p_j + t$. The first order conditions hold when both buyers price above the reservation price.

Proof of Lemma 3.

Case 1: $|p_i - p_j| \leqslant t$. We see from Equation (P-1) in Appendix B.1.1 that the profit maximization problem via the choice of pricing can be equivalently solved by maximizing profit via the choice of safety stock,

$$\max_{s_i(\lambda) \geqslant \frac{\lambda}{t}(\underline{p} - p_j) + 1 - Q_i} \pi_i(s_i(\lambda), p_j, \lambda),$$

and taking $p_i(p_j,\lambda) = p_j + \frac{t}{\lambda}(Q_i - 1) + \frac{t}{\lambda}\cdot s_i(\lambda)$.

The partial derivative $\frac{\partial \pi_i(s_i(\lambda),p_j,\lambda)}{\partial s_i(\lambda)}$ is given in Equation (P-FOC) in Appendix B.1.1 and is continuous in $s_i(\lambda)$. If $\pi_i(s_i(\lambda),p_j,\lambda)$ is piecewise concave in s_i, then the expected profit is strictly concave in $s_i(\lambda)$ so the FOC gives the global optimal. We verify that $\pi_i(s_i(\lambda),p_j,\lambda)$ is piecewise concave by taking the second derivative,

$$\frac{\partial^2 \pi_i(s_i(\lambda),p_j,\lambda)}{\partial s_i(\lambda)^2} = \begin{cases} -2\cdot\frac{t}{\lambda} - (\gamma+h)f(s_i(\lambda);\sigma) < 0 & \text{if } F(s_i(\lambda);\sigma)\in(0,1), \\ -2\cdot\frac{t}{\lambda} < 0 & \text{if } F(s_i(\lambda);\sigma) = 0 \,\text{or}\, 1. \end{cases}$$

Let $s_i^{FOC}(\lambda)$ satisfy $\left.\frac{\partial \pi_i(s_i(\lambda),p_j,\lambda)}{\partial s_i(\lambda)}\right|_{s_i(\lambda)=s_i^{FOC}(\lambda)} = 0$. This is equivalent to $s_i^{FOC}(\lambda)$ satisfying Equation (FOC-2). Since the LHS of (FOC-2) is monotonically increasing in $s_i^{FOC}(\lambda)$, it has at most one solution.

Let $p_i^{FOC}(p_j,\lambda)$ be defined as in Equation (FOC-1). The piece-wise concavity of $\pi_i(s_i(\lambda),p_j,\lambda)$, and the fact that $p_i^{FOC}(p_j,\lambda)$ is increasing $s_i^{FOC}(\lambda)$, imply that $\max\{\underline{p}, p_i^{FOC}(\lambda)\}$ is the unique optimizer for the expected profit expression given in Equation (P-1) under the constraint $p_i \geqslant \underline{p}$.

Case 2: $|p_i - p_j| > t$. The expected profit and the best price response when $|p_i - p_j| \geqslant t$ are given in Appendix B.1.2. It is clear that the best response is always capped at $p_j + t$. We thus define

$$\tilde{p}_i^{BR}(p_j,\lambda) = \min\left\{\max\left\{\underline{p}, p_i^{FOC}(p_j,\lambda)\right\}, p_j + t\right\},$$

and recognize that $\tilde{p}_i^{BR}(p_j,\lambda)$ can be sub-optimal in two cases.

- When $p_j - t \geqslant \tilde{p}_i^{BR}(p_j,\lambda) > \underline{p}$, lowering the price from $p_i = \tilde{p}_i^{BR}(p_j,\lambda)$ to $p_i = \underline{p}$ always strictly increases buyer i's profit for $\lambda < 1$ and has no impact on buyer i's profit for $\lambda = 1$.
- When $\tilde{p}_i^{BR}(p_j,\lambda) > p_j - t > \underline{p}$, buyer i's expected profit first decreases when p_i is lowered from $\tilde{p}_i^{BR}(p_j,\lambda)$ to $p_j - t$, and then increases when p_i is lowered further from $p_j - t$ to $\underline{p}$. As a result, the optimal price response is determined by $\arg\max_{p_i\in\{\underline{p},\tilde{p}_i^{BR}\}} \pi_i(p_i,p_j,\lambda)$.

Putting all the cases together, we have the result in Lemma 3. □

The next lemma recasts the first order conditions in terms of conditions on the average price and price dispersion.

Lemma 4. *If $p_H^*(\lambda) > \underline{p}$, $p_L^*(\lambda) > \underline{p}$ is a non-cooperative pure-strategy Bayes-Nash equilibrium such that $|p_H^*(\lambda) - p_L^*(\lambda)| < t$, then $p_H^*(\lambda)$, $p_L^*(\lambda)$ must satisfy the simultaneous equations in* (FOC-3), *i.e.,*

$$\frac{p_H^*(\lambda) + p_L^*(\lambda)}{2} = 1 + \frac{\gamma - h}{2} - \frac{t}{\lambda}, \quad p_H^*(\lambda) - p_L^*(\lambda) = \frac{\gamma+h}{3}(1 - 2F(s_H(p_H^*,p_L^*,\lambda);\sigma)),$$

and the solution is unique. Moreover, if $\gamma + h > 0$ and $\Delta > 0$ then $0 < \frac{\lambda}{t}(p_H^*(\lambda) - p_L^*(\lambda)) < \Delta$.

Proof of Lemma 4. If $p_H^* > \underline{p}$, $p_L^* > \underline{p}$ is an equilibrium such that $|p_H^*(\lambda) - p_L^*(\lambda)| < t$, both must satisfy the first order condition in Equation (FOC-2), i.e.,

$$2s_H^{FOC}(p_H^*,p_L^*,\lambda)\cdot\frac{t}{\lambda} + (\gamma+h)F(s_H^{FOC}(p_H^*,p_L^*,\lambda);\sigma) \;=\; 1 - p_L^* + \gamma + \frac{t}{\lambda}(1 - 2Q_H) \tag{16}$$

$$2s_L^{FOC}(p_L^*,p_H^*,\lambda)\cdot\frac{t}{\lambda} + (\gamma+h)F(s_L^{FOC}(p_L^*,p_H^*,\lambda);\sigma) \;=\; 1 - p_H^* + \gamma + \frac{t}{\lambda}(1 - 2Q_L). \tag{17}$$

Adding Equations (16) and (17) yields

$$\frac{p_H^*(\lambda) + p_L^*(\lambda)}{2} = 1 + \frac{\gamma - h}{2} - \frac{t}{\lambda},$$

where we used the facts that $s_H^{FOC}(\lambda)+s_L^{FOC}(\lambda)=0$, $F(s_L^{FOC}(\lambda);\sigma)+F(s_H^{FOC}(\lambda);\sigma)=1$, and $Q_H+Q_L=2$. Subtracting Equation (17) from (16) yields

$$p_H^*(\lambda)-p_L^*(\lambda)=\frac{\gamma+h}{3}(1-2F(s_H(p_H^*,p_L^*,\lambda);\sigma))\equiv\frac{\gamma+h}{3}(1-2F(\frac{\lambda}{t}(p_H^*(\lambda)-p_L^*(\lambda))-\Delta;\sigma)). \tag{18}$$

Note that the RHS is non-increasing in $p_H(\lambda)-p_L(\lambda)$ since $s_H(\lambda)$ is strictly increasing in $p_H(\lambda)-p_L(\lambda)$ and $F(s_H(\lambda);\sigma)$ is non-decreasing in $s_H(\lambda)$. Hence the solution to the two equations is unique as required.

If $\Delta>0,\gamma+h>0$, then Equation (18) shows that $p_H^*(\lambda)-p_L^*(\lambda)>0$ (otherwise, $RHS>0$ and $LHS\leqslant 0$ lead to a contradiction). Similarly, we must have $s_H(p_H^*,p_L^*,\lambda)<0$ (otherwise, $RHS<0$ and $LHS\geqslant 0$ lead to a contradiction). Since $s_H^*(\lambda)=\frac{\lambda}{t}p_\Delta^*(\lambda)-\Delta$, this shows that $\frac{\lambda}{t}p_\Delta^*(\lambda)<\Delta$. □

C.3. Technical Lemmas

In this section, we provide a few technical lemmas which facilitate a full characterization of the non-cooperative Bayes-Nash equilibrium under any transparency level $\lambda\in(0,1]$. The lemmas in this section use the definitions of $p_i^{FOC}(p_j,\lambda)$ and $s_i^{FOC}(p_i,p_j,\lambda)$ in the first order conditions (FOC-1) and (FOC-2).

The first lemma shows that, when best responding to the same price, the high-demand buyer prices higher than the low-demand buyer.

Lemma 5. *For all $\lambda\in(0,1]$ and prices $p\geqslant \underline{p}$ it holds that $p_H^{FOC}(p,\lambda)\geqslant p_L^{FOC}(p,\lambda)$.*

Proof. Suppose for the sake of contradiction that there exists some λ,p such that $p_H^{FOC}(p,\lambda)<p_L^{FOC}(p,\lambda)$. The first order condition in Equation (FOC-1) implies that

$$s_H^{FOC}(p,\lambda)-s_L^{FOC}(p,\lambda)=\frac{\lambda}{t}\left(p_H^{FOC}(p,\lambda)-p_L^{FOC}(p,\lambda)\right)-2\Delta$$

for any price p, so if $p_H^{FOC}(p,\lambda)<p_L^{FOC}(p,\lambda)$ then $s_H^{FOC}(p,\lambda)-s_L^{FOC}(p,\lambda)<-2\Delta<0$.

Similarly, the first order condition in Equation (FOC-2) gives

$$\begin{aligned}2s_H^{FOC}(p,\lambda)\cdot\frac{t}{\lambda}+(\gamma+h)F(s_H^{FOC}(p,\lambda);\sigma)&=1-p+\gamma+\frac{t}{\lambda}(1-2Q_H),\\ 2s_L^{FOC}(p,\lambda)\cdot\frac{t}{\lambda}+(\gamma+h)F(s_L^{FOC}(p,\lambda);\sigma)&=1-p+\gamma+\frac{t}{\lambda}(1-2Q_L).\end{aligned}$$

Subtracting the two equations yields

$$\begin{aligned}-4\frac{t}{\lambda}\Delta= &\quad 2[s_H^{FOC}(p,\lambda)-s_L^{FOC}(p,\lambda)]\cdot\frac{t}{\lambda}+(\gamma+h)[F(s_H^{FOC}(p,\lambda);\sigma)-F(s_L^{FOC}(p,\lambda);\sigma)]\\ < &\quad -4\frac{t}{\lambda}\Delta+(\gamma+h)[F(s_H^{FOC}(p,\lambda);\sigma)-F(s_L^{FOC}(p,\lambda);\sigma)] \quad \textit{(by assumption on p)}\\ \leqslant &\quad -4\frac{t}{\lambda}\Delta \quad \textit{(since } s_H^{FOC}(p,\lambda)<s_L^{FOC}(p,\lambda) \textit{ and } F(s;\sigma) \textit{ is weakly increasing in s)},\end{aligned}$$

which gives the required contradiction. □

The next lemma states that $p_H^{FOC}(p_L,\lambda)$ is strictly increasing in p_L and characterizes the range of the derivative.

Lemma 6. *For any given p_L, $\frac{\partial p_H^{FOC}(p_L,\lambda)}{\partial p_L}\in\left[\frac{1}{2},1\right)$ for all $\lambda\in(0,1]$.*

Proof. Differentiating the first order conditions in Equations (FOC-1) and (FOC-2) for $i=H$ with respect to p_L:

$$2\frac{\partial s_H^{FOC}(p_L,\lambda)}{\partial p_L}\cdot\frac{t}{\lambda}+(\gamma+h)\frac{\partial s_H^{FOC}(p_L,\lambda)}{\partial p_L}f(s_H^{FOC}(p_L,\lambda);\sigma)=-1$$

$$\Rightarrow \qquad \frac{\partial s_H^{FOC}(p_L,\lambda)}{\partial p_L}\cdot\frac{t}{\lambda} = \frac{-1}{2+(\gamma+h)\frac{\lambda}{t}f(s_H^{FOC}(p_L,\lambda);\sigma)} \in \left[-\frac{1}{2},0\right)$$

$$\Rightarrow \qquad \frac{\partial p_H^{FOC}(p_L,\lambda)}{\partial p_L} = 1+\frac{\partial s_H^{FOC}(p_L,\lambda)}{\partial p_L}\cdot\frac{t}{\lambda}\in\left[\frac{1}{2},1\right) \quad \text{as required.} \quad \square$$

The next lemma establishes the necessary and sufficient condition for $p_H^{FOC}(\underline{p},\lambda)\leqslant \underline{p}$.

Lemma 7. $\left.\frac{\partial\pi_H(s_H(\lambda),p_L,\lambda)}{\partial p_H(\lambda)}\right|_{p_H=p_L=\underline{p}}\leqslant 0$ *and* $p_H^{FOC}(\underline{p},\lambda)\leqslant\underline{p}$ *if and only if* $(\gamma+h)F(\Delta;\sigma)-\frac{t}{\lambda}\leqslant\underline{p}+h-1$.

Proof. Using the first order condition in Equation (P-FOC) in Appendix B.1.1 and the fact that $p_H = p_L=\underline{p}$ implies that $s_H(\lambda)=1-Q_H=-\Delta$, we have

$$\begin{aligned}\left.\frac{\partial\pi_H(s_H(\lambda),p_L,\lambda)}{\partial s_H(\lambda)}\right|_{p_H=p_L=\underline{p}} &= 2\Delta\cdot\frac{t}{\lambda}+1-\underline{p}-\frac{t}{\lambda}(1+2\Delta)+\gamma[1-F(-\Delta;\sigma)]-hF(-\Delta;\sigma)\\ &= 1-\underline{p}-h-\frac{t}{\lambda}+(\gamma+h)F(\Delta;\sigma).\end{aligned}$$

Thus, $\left.\frac{\partial\pi_H(s_H(\lambda),p_L,\lambda)]}{\partial p_H(\lambda)}\right|_{p_H=p_L=\underline{p}}\leqslant 0 \iff (\gamma+h)F(\Delta;\sigma)-\frac{t}{\lambda}\leqslant\underline{p}+h-1$. Meanwhile, piece-wise concavity of $\pi_H(s_H(\lambda),p_L,\lambda)$ implies that $\left.\frac{\partial\pi_H(s_H(\lambda),p_L,\lambda)]}{\partial p_H(\lambda)}\right|_{p_H=p_L=\underline{p}}\leqslant 0 \iff p_H^{FOC}(\underline{p},\lambda)\leqslant\underline{p}$. $\square$

The next lemma rules out equilibria where only the high-demand buyer prices at the reservation price.

Lemma 8. *There is no Bayes-Nash equilibrium where* $p_H^*(\lambda)=\underline{p}$, $p_L^*(\lambda)>\underline{p}$ *for any* $\lambda\in(0,1]$.

Proof. Suppose for the sake of contradiction there is a pure strategy equilibrium with $p_L^*(\lambda)>\underline{p}$ and $p_H^*(\lambda)=\underline{p}$. For p_L^* to be optimal for buyer L, we must have $p_L^*(\lambda)\leqslant\underline{p}+t$. The conditions for equilibrium are

$$p_L^*+\frac{t}{\lambda}(\Delta+s_H^{FOC}(p_L^*,\lambda))\leqslant\underline{p} \quad \text{(optimal response } p_H^*=\underline{p})$$

$$\underline{p}+\frac{t}{\lambda}(-\Delta+s_L^{FOC}(\underline{p},\lambda))>\underline{p} \quad \text{(optimal response } p_L^*>\underline{p}).$$

Subtracting the two equations gives

$$\begin{aligned}0 &\geqslant (p_L^*-\underline{p})+\frac{t}{\lambda}(2\Delta+s_H^{FOC}(p_L^*,\lambda)-s_L^{FOC}(\underline{p},\lambda))\\ &= p_H^{FOC}(p_L^*,\lambda)-p_L^{FOC}(\underline{p},\lambda) \qquad \text{(by Equation (FOC-1))}\\ &> p_H^{FOC}(\underline{p},\lambda)-p_L^{FOC}(\underline{p},\lambda) \qquad \text{(by Lemma 6, as } p_L^*>\underline{p})\\ &\geqslant 0 \qquad \text{(by Lemma 5).}\end{aligned}$$

This gives the required contradiction. $\square$

C.4. Characterizing $\hat{t}$

For $p_H^*(\lambda),p_L^*(\lambda)$ in our full characterization in Section C.1 to be a Bayes-Nash equilibrium, we require two additional conditions on the transportation cost, t.

First, we need t to be sufficiently large such that $p_H^*(\lambda)-p_L^*(\lambda)\leqslant t$ for all $\lambda\in(0,1]$. Otherwise, buyer H can always improve his expected profit by lowering his price from $p_H^*(\lambda)$ to $p_L^*(\lambda)+t$ (see Appendix B.1.2) and the market is not in equilibrium.

Let $\hat{t}_1$ be the minimum value of t such that for all $t\geqslant\hat{t}_1$,

$$p_H^{FOC}(\underline{p},\lambda)\leqslant\underline{p}+t \quad \forall \quad \lambda\in(0,1].$$

Given Lemma 6, this provides a sufficient condition for $p^*_H(\lambda) - p^*_L(\lambda) \leqslant t$ to hold for all $\lambda \in (0,1]$ whenever $t \geqslant \hat{t}_1$. We know that such a $\hat{t}_1$ always exists because, from Lemma 7, $p_H^{FOC}(\underline{p},\lambda) \leqslant \underline{p}$ for sufficiently large t, and $p_H^{FOC}(\underline{p},\lambda)$ is continuous in t.

Second, whenever $p^*_i(\lambda) > p^*_j(\lambda) - t > \underline{p}$, we need to make sure that it is never optimal for buyer i to deviate from $p^*_i(\lambda)$ to price at $\underline{p}$ instead (see Lemma 3 scenario (2)). Such deviation becomes profitable if $p^*_i(\lambda)$ is so high that buyer i is better off charging the monopolistic price and buying only from its loyal base. Lemma 9 states that there exists a threshold $\hat{t}$ on t such that if $t \geqslant \hat{t}$ then for all $\lambda \in (0,1]$ the deviation to $\underline{p}$ is never optimal.

Lemma 9. *There exists $\hat{t} \geqslant \hat{t}_1$ such that for all $t \geqslant \hat{t}$, whenever both buyers price above $\underline{p}$ and $p^*_H(\lambda) \geqslant p^*_L(\lambda) > \underline{p} + t$, it holds that $\pi_i(p^*_i, p^*_j, \lambda) \geqslant \pi_i(\underline{p}, p^*_j, \lambda)$ for $i \in \{H, L\}$ for all $\lambda \in (0,1]$.*

Proof. We focus on $t \geqslant \hat{t}_1$ such that $p^*_H(\lambda) - p^*_L(\lambda) \leqslant t$. For this proof, we consider any fixed value of λ and write all relevant quantities as functions of t instead of λ.

Since $p^*_H(t) \geqslant p^*_L(t) > \underline{p} + t$, the expected profits of buyer i under $p^*_H(t), p^*_L(t)$, $\pi_i(p^*_i(t), p^*_j(t), t)$ for $i \in \{H, L\}$, are given by Equation (P-1). The expected profit when buyer i deviates from p^*_i to price at $\underline{p}$, $\pi_i(\underline{p}, p^*_j(t), t)$, is given by Equation (P-2). We will prove the lemma by showing that there are constants c_H, c_L such that for any fixed λ,

$$\frac{d\pi_i(p^*_i(t), p^*_j(t), t)}{dt} > c_i > 0, \quad \frac{d\pi_i(\underline{p}, p^*_j(t), t)}{dt} = 0 \quad \text{for } i \in \{H, L\}.$$

Using Equation (P-2), it is straightforward to verify that $\frac{d\pi_i(\underline{p}, p^*_j(t), t)}{dt} = 0$. We now show $\frac{d\pi_i(p^*_i(t), p^*_j(t), t)}{dt} > 0$ using Equation (P-1). First, note that $\frac{\partial \pi_i(s^*_i(t), p^*_j(t), t)}{\partial s^*_i(t)} = 0$ since $p^*_H(t) = \tilde{p}_H(t), p^*_L(t) = \tilde{p}_L(t)$ satisfy the first order conditions. Thus, differentiating both sides of Equation (P-1) gives

$$\begin{aligned}\frac{d\pi_i(s^*_i(t), p^*_j(t), t)}{dt} &= \frac{\partial \pi_i(s^*_i(t), p^*_j(t), t)}{\partial p^*_j(t)} \cdot \frac{dp^*_j(t)}{dt} + \frac{\partial \pi_i(s^*_i(t), p^*_j(t), t)}{\partial t} \\ &= (Q_i + s^*_i(t)) \left[-\frac{dp^*_j(t)}{dt} - \frac{1}{\lambda}(s_i(t) + Q_i - 1) \right].\end{aligned}$$

Letting $\bar{p}^*(t) = \frac{1}{2}(p^*_H(t) + p^*_L(t))$ and $p^*_\Delta(t) = p^*_H(t) - p^*_L(t)$ denote the average price and price dispersion, we have

$$\frac{dp^*_H(t)}{dt} = \frac{\partial \bar{p}^*(t)}{dt} + \frac{1}{2}\frac{dp^*_\Delta(t)}{dt}, \quad \frac{dp^*_L(t)}{dt} = \frac{\partial \bar{p}^*(t)}{dt} - \frac{1}{2}\frac{dp^*_\Delta(t)}{dt},$$

and, by differentiating both equations in Equation (FOC-3) with respect to t, we have

$$\frac{\partial \bar{p}^*(t)}{dt} = -\frac{1}{\lambda}, \quad \frac{dp^*_\Delta(t)}{dt} = \frac{p^*_\Delta(t)}{t} \cdot \frac{\frac{\gamma+h}{3}\frac{2\lambda}{t} f(s^*_H(t))}{1 + \frac{\gamma+h}{3}\frac{2\lambda}{t} f(s^*_H(t))} \in \left(0, \frac{p^*_\Delta(t)}{t}\right).$$

Moreover, multiplying both sides of the second equation in Equation (FOC-3) by $\frac{\lambda}{t}$ and differentiating with respect to t yields

$$\frac{d(\frac{\lambda}{t}p^*_\Delta(t))}{dt} = -\frac{\frac{\lambda}{t^2}\frac{\gamma+h}{3}[1 - 2F(\frac{\lambda}{t}p^*_\Delta(t) - \Delta)]}{1 + 2\frac{\lambda}{t}\frac{\gamma+h}{3} f(\frac{\lambda}{t}p^*_\Delta(t))} \leqslant 0,$$

where we used $\frac{\lambda}{t}p^*_\Delta(t) \leqslant \Delta$ (Lemma 4) and $F(\frac{\lambda}{t}p^*_\Delta(t) - \Delta) \leqslant F(0) = \frac{1}{2}$ since $f(\cdot)$ is symmetric around 0.

For buyer H,

$$\begin{aligned}\frac{d\pi_H(s^*_H(t), p^*_L(t), t)}{dt} &= (1 + \Delta + s^*_H(t)) \left[-\left(-\frac{1}{\lambda} - \frac{1}{2}\frac{dp^*_\Delta(t)}{dt}\right) - \frac{1}{t}p^*_\Delta(t) \right] \\ &= \frac{1}{\lambda}\left(1 + \frac{\lambda}{t}p^*_\Delta(t)\right)\left[1 - \frac{\lambda}{t}p^*_\Delta(t) \cdot \left(1 - \frac{\frac{\gamma+h}{3}\frac{\lambda}{t} f(s^*_H(t))}{1 + \frac{\gamma+h}{3}\frac{2\lambda}{t} f(s^*_H(t))}\right)\right] \geqslant \frac{1}{\lambda}\left[1 - \left(\frac{\lambda}{t}p^*_\Delta(t)\right)^2\right] > 0.\end{aligned}$$

The inequalities follow from $0 \leqslant \frac{\lambda}{t} p^*_\Delta(t) < \Delta \leqslant 1$ if $\Delta > 0$ (Lemma 4). Moreover, the lower bound $\frac{1}{\lambda}\left[1 - \left(\frac{\lambda}{t} p^*_\Delta(t)\right)^2\right]$ is increasing in t since $\frac{d(\frac{\lambda}{t} p^*_\Delta(t))}{dt} \leqslant 0$. Thus $\frac{d\pi_H(s^*_H(t), p^*_L(t), t)}{dt}$ does not converge to 0.

For buyer L,

$$\begin{aligned}\frac{d\pi_L(s^*_L(t), p^*_H(t), t)}{dt} &= (1 - \Delta + s^*_L(t)) \left[- \left(-\frac{1}{\lambda} + \frac{1}{2}\frac{dp^*_\Delta(t)}{dt}\right) + \frac{1}{t} p^*_\Delta(t)\right] \\ &= \frac{1}{\lambda}(1 - \frac{\lambda}{t} p^*_\Delta(t)) \left[1 + \frac{\lambda}{t} p^*_\Delta(t) \left(1 - \frac{\frac{\gamma+h}{3}\frac{\lambda}{t} f(s^*_H(t))}{1 + \frac{\gamma+h}{3}\frac{2\lambda}{t} f(s^*_H(t))}\right)\right] \geqslant \frac{1}{\lambda}\left(1 - \frac{\lambda}{t} p^*_\Delta(t)\right) > 0.\end{aligned}$$

The inequalities again follow from $0 \leqslant \frac{\lambda}{t} p^*_\Delta(t) < \Delta \leqslant 1$ if $\Delta > 0$. Similarly, the lower bound is increasing in t and thus $\frac{d\pi_L(s^*_L(t), p^*_H(t), t)}{dt}$ does not converge to 0. Therefore, there always exists $\hat{t} \geqslant \hat{t}_1$ such that for all $t \geqslant \hat{t}$, $\pi_i(p^*_i, p^*_j, \lambda) \geqslant \pi_i(\underline{p}, p^*_j, \lambda)$ for $i \in \{H, L\}$ for any fixed λ. □

We thus use $\hat{t}$ to denote the minimum value of t such that $p^{FOC}_H(\underline{p}, \lambda) \leqslant \underline{p} + t$ and $\pi_i(p^*_i, p^*_j, \lambda) \geqslant \pi_i(\underline{p}, p^*_j, \lambda)$ for $i \in \{H, L\}$ for all λ such that $p^*_H(\lambda) \geqslant p^*_L(\lambda) > \underline{p} + t$, where $p^*_H(\lambda), p^*_L(\lambda)$ are given by the above characterization.

C.5. Proof of Full Characterization and Theorem 1

We prove Theorem 1 by proving the validity and uniqueness of the full characterization stated in Appendix C.1 for all $t \geqslant \hat{t}$, where $\hat{t}$ is characterized in Appendix C.4. This proves the existence of a unique Bayes-Nash equilibrium, and shows the equilibrium prices satisfy conditions (4)-(7) in Theorem 1.

Case 1. We characterize when at equilibrium both buyers price at $\underline{p}$.

Using Lemma 7, it suffices to show that $p^*_H(\lambda) = p^*_L(\lambda) = \underline{p}$ is a Bayes-Nash equilibrium if and only if $\left.\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial p_H(\lambda)}\right|_{p_H = p_L = \underline{p}} \leqslant 0$, and that if $p^*_H(\lambda) = p^*_L(\lambda) = \underline{p}$ is a Bayes-Nash equilibrium then it is the unique Nash equilibrium.

We first show that if $p_H = p_L = \underline{p}$ is a Bayes-Nash equilibrium then $\left.\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)}\right|_{p_H = p_L = \underline{p}} \leqslant 0$. This follows from the fact that $p_H = \underline{p}$ is a best response for the high-demand buyer, and so $\left.\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)}\right|_{p_H = p_L = \underline{p}} \leqslant 0$.

Next, we show that if $\left.\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)}\right|_{p_H = p_L = \underline{p}} \leqslant 0$ then $p_H = p_L = \underline{p}$ is a Bayes-Nash equilibrium. If $\left.\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)}\right|_{p_H = p_L = \underline{p}} \leqslant 0$ then $p^{FOC}_H(\underline{p}, \lambda) \leqslant \underline{p}$. Lemma 5 implies that $p^{FOC}_L(\underline{p}, \lambda) \leqslant p^{FOC}_H(\underline{p}, \lambda) \leqslant \underline{p}$.

So far, we have shown that $p^*_H(\lambda) = p^*_L(\lambda) = \underline{p}$ is a Bayes-Nash equilibrium if and only if $(\gamma + h)F(\Delta; \sigma) - \frac{t}{\lambda} \leqslant \underline{p} + h - 1$. Lastly, we show that if $p^*_H(\lambda) = p^*_L(\lambda) = \underline{p}$ is a Bayes-Nash equilibrium, it is the unique Bayes-Nash equilibrium. It is obvious that $p_L(\lambda) = \underline{p}, p_H(\lambda) > \underline{p}$ is not an equilibrium because $\left.\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial p_H(\lambda)}\right|_{p_H = p_L = \underline{p}} \leqslant 0$ and $\pi_H(s_H(\lambda), p_L, \lambda)$ is strictly concave in p_H. Given Lemma 8, it suffices to show that there is no Bayes-Nash equilibrium with $p_H > \underline{p}, p_L > \underline{p}$.

Suppose for the sake of contradiction there is also a Bayes-Nash equilibrium (p_H, p_L) with $p_H > \underline{p}, p_L > \underline{p}$. Then the average price in this equilibrium satisfies

$$\underline{p} < \frac{p_H + p_L}{2} = 1 + \frac{\gamma - h}{2} - \frac{t}{\lambda} \leqslant 1 - h + (\gamma + h)F(\Delta; \sigma) - \frac{t}{\lambda} \leqslant \underline{p},$$

where the first inequality holds since $p_H, p_L > \underline{p}$, the equality holds by Lemma 4, the second inequality holds since $F(\Delta; \sigma) \geqslant \frac{1}{2}$ (as $\Delta \geqslant 0$ and $F(\cdot; \sigma)$ is symmetric around 0), and the last inequality holds since there is Bayes-Nash equilibrium with $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$. This gives the required contradiction and completes the proof of Case 1.

Case 2. We characterize when at equilibrium both buyers price above $\underline{p}$. Suppose $\tilde{p}_L(\lambda) > \underline{p}$. It is straightforward to verify that $\tilde{p}_H, \tilde{p}_L > \underline{p}$ satisfies the FOC, and hence is a Bayes-Nash equilibrium. We first show that in this case, $\tilde{p}_H \geqslant \tilde{p}_L$. Note that by definition of $\tilde{p}_H$ and $\tilde{p}_L$

$$\tilde{p}_H - \tilde{p}_L = \frac{\gamma + h}{3}(1 - 2F(s_H(\tilde{p}_H, \tilde{p}_L, \lambda); \sigma)), \tag{19}$$

so it suffices to show that $F(s_H(\tilde{p}_H, \tilde{p}_L, \lambda); \sigma) \leqslant 1/2$, or equivalently $s_H(\tilde{p}_H, \tilde{p}_L, \lambda) \leqslant 0$ (as $\Delta \geqslant 0$ and $F(\cdot; \sigma)$ is symmetric around 0). Note that $\tilde{p}_H, \tilde{p}_L$ satisfy $\frac{\lambda}{t}(\tilde{p}_H - \tilde{p}_L) = s_H(\tilde{p}_H, \tilde{p}_L, \lambda) + \Delta$ by the definition of s_H. Plugging this into Equation (19) yields

$$-s_H(\tilde{p}_H, \tilde{p}_L, \lambda) + \frac{\lambda}{t}\frac{\gamma + h}{3}[1 - 2F(s_H(\tilde{p}_H, \tilde{p}_L; \sigma), \lambda)] = \Delta \geqslant 0$$

If $s_H(\tilde{p}_H, \tilde{p}_L, \lambda) > 0$, then both terms on the LHS are negative. This implies that $s_H(\tilde{p}_H, \tilde{p}_L, \lambda) \leqslant 0$.

We now show that if $\tilde{p}_L > \underline{p}$ then $(\tilde{p}_H, \tilde{p}_L)$ is the *unique* Bayes-Nash equilibrium. The only other possibility of being at equilibrium without satisfying the FOC is that any buyer i pricing at the $\underline{p}$ satisfies $\frac{\partial \pi_i(s_i(\lambda), p_j, \lambda)}{\partial s_i(\lambda)} \leqslant 0$. Given Lemma 8 and Case 1, it suffices to show that whenever $p_L = \underline{p}$ and $p_H = p_H^{FOC}(\underline{p}, \lambda)$, $\frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)} > 0$ and therefore the market is not at equilibrium.

Using the result of Lemma 6, we have

$$\frac{\partial p_H^{FOC}(p_L, \lambda)}{\partial p_L} > 0 \quad \forall \lambda \in (0, 1). \tag{20}$$

Using the FOC in Equation (P-FOC), we have

$$\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)]}{\partial s_L(\lambda)} = 2 - 2\frac{t}{\lambda} + \gamma - h - (p_L + p_H). \tag{21}$$

When $p_H + p_L = \tilde{p}_H(\lambda) + \tilde{p}_L(\lambda) \equiv p_H^{FOC}(\tilde{p}_L(\lambda), \lambda) + \tilde{p}_L(\lambda)$, Lemma 4 implies that

$$\left(\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)}\right)\bigg|_{\tilde{p}_H, \tilde{p}_L} = 0.$$

When $p_L = \underline{p} < \tilde{p}_L$, Equation (20) implies that $p_H = p_H^{FOC}(\underline{p}, \lambda) < \tilde{p}_H(\lambda)$, and thus $p_H + p_L < \tilde{p}_H(\lambda) + \tilde{p}_L(\lambda)$. This means when $p_L = \underline{p}$ and $p_H = p_H^{FOC}(\underline{p}, \lambda)$,

$$\frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)} = \frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)} > 0 \quad \text{as required,}$$

where the equality holds since p_H satisfies the FOC, and the inequality follows from Equation (21) since $p_H + p_L < \tilde{p}_H(\lambda) + \tilde{p}_L(\lambda)$ and $\left(\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)}\right)\Big|_{\tilde{p}_H, \tilde{p}_L} = 0$.

Case 3. We characterize when exactly one buyer prices above reservation price. The last remaining case satisfies $(\gamma + h)F(\Delta; \sigma) - \frac{t}{\lambda} > \underline{p} + h - 1$ and $\tilde{p}_L \leqslant \underline{p}$. The results in Cases 1 and 2 and Lemma 8 show the only possible equilibrium is $p_L^*(\lambda) = \underline{p}$ and $p_H^*(\lambda) = p_H^{FOC}(\underline{p}, \lambda)$. We show that this is indeed an equilibrium, which is equivalent to showing $\frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)} \leqslant 0$.

We proceed as in the proof of Case 2. When $p_H + p_L = \tilde{p}_H(\lambda) + \tilde{p}_L(\lambda) \equiv p_H^{FOC}(\tilde{p}_L(\lambda), \lambda) + \tilde{p}_L(\lambda)$, Lemma 4 implies that $\left(\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)}\right)\Big|_{\tilde{p}_H, \tilde{p}_L} = 0$.

When $p_L = \underline{p} \geqslant \tilde{p}_L$, Equation (20) implies that $p_H = p_H^{FOC}(\underline{p}, \lambda) \geqslant \tilde{p}_H(\lambda)$, and thus $p_H + p_L \geqslant \tilde{p}_H(\lambda) + \tilde{p}_L(\lambda)$. This means when $p_L = \underline{p}$ and $p_H = p_H^{FOC}(\underline{p}, \lambda)$,

$$\frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)} = \frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)} \leqslant 0 \quad \text{as required,}$$

where the equality holds since p_H satisfies the FOC, and the inequality follows from Equation (21) since $p_H + p_L \geqslant \tilde{p}_H(\lambda) + \tilde{p}_L(\lambda)$ and $\left(\frac{\partial \pi_H(s_H(\lambda), p_L, \lambda)}{\partial s_H(\lambda)} + \frac{\partial \pi_L(s_L(\lambda), p_H, \lambda)}{\partial s_L(\lambda)}\right)\Big|_{\tilde{p}_H, \tilde{p}_L} = 0$.

This concludes the proof of Theorem 1.

Appendix D: Remaining Proofs

This section present proofs for all remaining analysis for the competitive and collusive equilibrium in Sections 3, 4 and 5.

We will make use of Lemma 10, which shows that in equilibrium, buyer H always buys less supply than his contractual demand in expectation, and buyer L always buys more in expectation.

Lemma 10. *For any $\lambda \in (0,1)$ such that $p_\Delta^*(\lambda) > 0$, it holds that $0 < \frac{\lambda}{t} p_\Delta^*(\lambda) < \Delta$, i.e.,*

$$1 < y(p_H^*, p_L^*, \lambda) < Q_H, \quad Q_L < y(p_L^*, p_H^*, \lambda) < 1.$$

Proof. For any λ such that $p_H^*(\lambda) > \underline{p}$ and $p_L^*(\lambda) > \underline{p}$, the result is already proven in Lemma 4. Now, consider any λ such that $p_H^*(\lambda) > \underline{p}$, $p_L^*(\lambda) = \underline{p}$. For such λ, $p_\Delta^* > 0$ is trivial. We show that $\frac{\lambda}{t} p_\Delta^*(\lambda) < \Delta$.

Let $\tilde{p}_H(\lambda), \tilde{p}_L(\lambda)$ denote the unique solution to the simultaneous equations

$$\frac{p_H(\lambda) + p_L(\lambda)}{2} = 1 + \frac{\gamma - h}{2} - \frac{t}{\lambda}, \quad p_H(\lambda) - p_L(\lambda) = \frac{\gamma + h}{3}(1 - 2F(s_H(p_H, p_L, \lambda); \sigma)).$$

Recall from the full characterization of equilibrium prices in Appendix C.1 that $\tilde{p}_L(\lambda) \leqslant \underline{p}$ for any λ such that $p_H^*(\lambda) > \underline{p}$, $p_L^*(\lambda) = \underline{p}$. Since $\frac{\partial\left(p_H^{FOC}(p_L, \lambda) - p_L\right)}{\partial p_L} \in (-\frac{1}{2}, 0)$ as a result of Lemma 6, this means

$$p_\Delta^* \leqslant \tilde{p}_H(\lambda) - \tilde{p}_L(\lambda) = \frac{\gamma + h}{3}(1 - 2F(s_H(\tilde{p}_H, \tilde{p}_L, \lambda); \sigma)).$$

This means that $s_H(p_H^*, p_L^*, \lambda) \leqslant s_H(\tilde{p}_H, \tilde{p}_L, \lambda)$. Meanwhile, $s_H(\tilde{p}_H, \tilde{p}_L, \lambda)$ satisfies

$$s_H(\tilde{p}_H, \tilde{p}_L, \lambda) = \frac{\lambda}{t} \frac{\gamma + h}{3}[1 - 2F(s_H(\tilde{p}_H, \tilde{p}_L, \lambda); \sigma)] - \Delta,$$

which implies $s_H(\tilde{p}_H, \tilde{p}_L, \lambda) < 0$ (otherwise, $RHS < 0$ and $LHS \leqslant 0$ lead to a contradiction). Hence $\frac{\lambda}{t} p_\Delta^*(\lambda) - \Delta = s_H(p_H^*, p_L^*, \lambda) \leqslant s_H(\tilde{p}_H, \tilde{p}_L, \lambda) < 0$, which shows $\frac{\lambda}{t} p_\Delta^*(\lambda) < \Delta$ as required. □

D.1. Proof of Proposition 1

1. First, we show that $\frac{dp_\Delta^*(\lambda)}{d\lambda} \leqslant 0$ for any λ where $p_H^*(\lambda), p_L^*(\lambda) > \underline{p}$. Differentiating both sides of Equation (7) in Theorem 1 with respect to λ,

$$\frac{dp_\Delta^*(\lambda)}{d\lambda} = -\frac{2(\gamma + h)}{3}\left(\frac{\partial F(\frac{\lambda}{t} p_\Delta^*(\lambda) - \Delta; \sigma)}{\partial \sigma} \frac{d\sigma}{d\lambda} + f\left(\frac{\lambda}{t} p_\Delta^*(\lambda) - \Delta; \sigma\right) \cdot \left(\frac{\lambda}{t} \frac{dp_\Delta^*(\lambda)}{d\lambda} + \frac{p_\Delta^*(\lambda)}{t}\right)\right)$$

Rearranging,

$$\frac{dp_\Delta^*(\lambda)}{d\lambda} = \frac{-\frac{2(\gamma+h)}{3} \frac{p_\Delta^*(\lambda)}{t} f\left(\frac{\lambda}{t} p_\Delta^*(\lambda) - \Delta; \sigma\right)}{\left[1 + \frac{2(\gamma+h)}{3} f\left(\frac{\lambda}{t} p_\Delta^*(\lambda) - \Delta; \sigma\right) \cdot \frac{\lambda}{t}\right]} \leqslant 0$$

In addition, it shows that the necessary and sufficient condition for $\frac{dp_\Delta^*(\lambda)}{d\lambda} < 0$ is $f(\frac{\lambda}{t} p_\Delta^*(\lambda) - \Delta; \sigma) > 0$, which is equivalent to $f(\Delta - \frac{\lambda}{t} p_\Delta^*(\lambda); \sigma) > 0$ since $f(\cdot; \sigma)$ is symmetric. The latter is equivalent to $0 <$

$F(\Delta - \frac{\lambda}{t}p^*_\Delta(\lambda);\sigma) < 1$, since the support of $f(\cdot;\sigma)$ is a bounded interval centered at 0 by Assumption 2. Note that $F(\Delta - \frac{\lambda}{t}p^*_\Delta(\lambda);\sigma) > 0$ is automatically guaranteed since $\Delta - \frac{\lambda}{t}p^*_\Delta(\lambda) > 0$ by Lemma 10.

2. Next, we show that $\frac{\partial p^*_\Delta(\lambda)}{\partial\sigma} \leqslant 0$ for $i \in \{H, L\}$ for any λ such that $p^*_\Delta(\lambda) > 0$. Since $s^*_H(\lambda) = \frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta$ this is equivalent to showing $\frac{\partial s^*_H(\lambda)}{\partial\sigma} \leqslant 0$.

(1) For any λ such that $p^*_H(\lambda) > p^*_L(\lambda) > \underline{p}$, Equation (7) in Theorem 1 gives that $p^*_\Delta(\lambda) = \frac{\gamma+h}{3}\left[1 - 2F\left(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma\right)\right]$. Differentiating both sides with respect to σ gives

$$\begin{aligned}\frac{\partial p^*_\Delta(\lambda)}{\partial\sigma} &= -\frac{2}{3}(\gamma+h)\left[f\left(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma\right)\frac{\partial p^*_\Delta(\lambda)}{\partial\sigma} + \frac{\partial F(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma)}{\partial\sigma}\right] \\ &= \frac{-\frac{2}{3}(\gamma+h)\frac{\partial F(\frac{\lambda}{t}p^*_\Delta(\lambda)-\Delta;\sigma)}{\partial\sigma}}{1 + \frac{2}{3}(\gamma+h)f\left(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma\right)},\end{aligned}$$

where $\frac{\partial F(\frac{\lambda}{t}p^*_\Delta(\lambda)-\Delta;\sigma)}{\partial\sigma} = -\frac{\frac{\lambda}{t}p^*_\Delta(\lambda)-\Delta}{\sigma}f(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma) \geqslant 0$ using Lemma 10. Thus, $\frac{\partial p^*_\Delta(\lambda)}{\partial\sigma} \leqslant 0$ with strict inequality if and only if $f(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma) > 0$.

(2) For any λ such that $p^*_H(\lambda) > \underline{p}$ and $p^*_L(\lambda) = \underline{p}$, Lemma 3 gives that $2s^*_H(\lambda)\cdot\frac{t}{\lambda} + (\gamma+h)F(s^*_H(\lambda);\sigma) = 1 - p_L + \gamma + \frac{t}{\lambda}(1 - 2Q_H)$. Differentiating both sides gives

$$\frac{\partial s^*_H(\lambda)}{\partial\sigma}\left(\frac{2t}{\lambda} + (\gamma+h)f(s^*_H(\lambda);\sigma)\right) - (\gamma+h)\frac{s^*_H(\lambda)}{\sigma}f(s^*_H(\lambda);\sigma) = 0.$$

Since $s^*_H(\lambda) < 0$ by Lemma 10, $\frac{\partial s^*_H(\lambda)}{\partial\sigma} \leqslant 0$ with strict inequality if and only if $f(\frac{\lambda}{t}p^*_\Delta(\lambda) - \Delta;\sigma) > 0$.

D.2. Proof of Lemma 2

From Equations (RP) and (L), the expected buyer profit is given by $\bar{\pi}(\mathbf{p},\lambda) = \bar{\Psi}(\mathbf{p},\lambda) - \bar{L}(\mathbf{p},\lambda)$, where $\bar{\Psi}(\mathbf{p},\lambda) = \frac{1}{2}\sum_{i\in\{H,L\}} y(p_i,p_j,\lambda)(1-p_i) = 1 - \bar{p} - \frac{\lambda}{2t}p^2_\Delta$ is maximized at $\mathbf{p} = \underline{\mathbf{p}}$, and $\bar{L}(\mathbf{p},\lambda) = (\gamma + h)\left[\int_{-\infty}^{-s(\mathbf{p},\lambda)} F(\epsilon;\sigma)d\epsilon + \frac{s(\mathbf{p},\lambda)}{2}\right]$. Observe that $\bar{L}(\mathbf{p},\lambda)$ depends on either $\mathbf{p}$ or λ only via $s(\mathbf{p},\lambda)$, and $\frac{\partial \bar{L}(\mathbf{p},\lambda)}{\partial s(\mathbf{p},\lambda)} = (\gamma+h)[\frac{1}{2} - F(-s(\mathbf{p},\lambda);\sigma)] > 0$, since $s(\mathbf{p},\lambda) > 0$ by Lemma 10.

D.3. Proof of Proposition 2

Proposition 2 assumes that supply uncertainty is a constant σ independent of λ. Here, we present the proof for a more general version of the proposition, where we relax the assumption of independence to $\frac{d\sigma}{d\lambda} \leqslant 0$, i.e., we allow σ to be either a constant or a function that is non-increasing in λ. We show that the result in Proposition 2 holds as long as $\frac{d\sigma}{d\lambda} \leqslant 0$, and the original proposition is a special case where $\frac{d\sigma}{d\lambda} = 0$.

First, we show that $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ for any $\lambda > \hat{\lambda}$.

(1) We first show that $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ for any λ such that $p^*_H(\lambda) > p^*_L(\lambda) > \underline{p}$. Using the expression for p^*_Δ in Equation (7), and noting that $s^*_H(\lambda) = \frac{\lambda}{t}(p^*_H(\lambda) - p^*_L(\lambda)) - \Delta$, we have $\frac{t}{\lambda}s^*_H(\lambda) = \frac{\gamma+h}{3}[1 - 2F(s^*_H(\lambda);\sigma)] - \frac{t}{\lambda}\Delta$. Differentiating both sides with respect to λ,

$$-\frac{t}{\lambda^2}s^*_H(\lambda) + \frac{t}{\lambda}\frac{ds^*_H(\lambda)}{d\lambda} = -2\frac{\gamma+h}{3}\left(\frac{\partial F(s^*_H(\lambda);\sigma)}{\partial\sigma}\frac{d\sigma}{d\lambda} + \frac{\partial F(s^*_H(\lambda);\sigma)}{\partial s^*_H(\lambda)}\frac{ds^*_H(\lambda)}{d\lambda}\right) + \frac{t}{\lambda^2}\Delta.$$

Rearranging and using that $\frac{\partial F(s^*_H(\lambda);\sigma)}{\partial\sigma} = -\frac{s^*_H(\lambda)}{\sigma}f(s^*_H(\lambda);\sigma) \geqslant 0$,

$$\underbrace{\left(2\frac{\gamma+h}{3}\frac{\partial F(s^*_H(\lambda);\sigma)}{\partial s^*_H(\lambda)} + \frac{t}{\lambda}\right)}_{>0}\frac{ds^*_H(\lambda)}{d\lambda} = \underbrace{\frac{t}{\lambda^2}\Delta}_{>0 \text{ if } \Delta>0} \underbrace{-2\frac{\gamma+h}{3}\frac{\partial F(s^*_H(\lambda);\sigma)}{\partial\sigma}\frac{d\sigma}{d\lambda}}_{\leqslant 0}.$$

The first term is always non-negative, the second term is positive if $\Delta > 0$, and the third term is non-negative if and only if $\frac{d\sigma}{d\lambda}$ is non-positive. Therefore $\frac{ds_H^*(\lambda)}{d\lambda} > 0$.

Since by Lemma 10 $s_H^*(\lambda) = -s(\mathbf{p}^*(\lambda), \lambda) < 0$, we have shown that $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ for $i \in \{H, L\}$.

(2) Next, we show that $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ for $i \in \{H, L\}$ for any λ such that $p_H^*(\lambda) > \underline{p}$ and $p_L^*(\lambda) = \underline{p}$. Recall from Lemma 10 and Equation (FOC-2) that $s_H^*(\lambda) < 0$ and satisfies the first order condition, $\frac{t}{\lambda} s_H^*(\lambda) = -\frac{\gamma+h}{2} F(s_H^*(\lambda); \sigma) + \frac{1}{2}(1 - \underline{p} + \gamma) - \frac{t}{\lambda}(1/2 + \Delta)$. Differentiating both sides with respect to λ,

$$-\frac{t}{\lambda^2} s_H^*(\lambda) + \frac{t}{\lambda} \frac{ds_H^*(\lambda)}{d\lambda} = -\frac{\gamma+h}{2}\left(\frac{\partial F(s_H^*(\lambda);\sigma)}{\partial \sigma}\frac{d\sigma}{d\lambda} + \frac{\partial F(s_H^*(\lambda);\sigma)}{\partial s_H^*(\lambda)}\frac{ds_H^*(\lambda)}{d\lambda}\right) + \frac{t}{\lambda^2}(\Delta + 1/2).$$

Rearranging and using that $\frac{\partial F(s_H^*(\lambda);\sigma)}{\partial \sigma} = -\frac{s_H^*(\lambda)}{\sigma} f(s_H^*(\lambda); \sigma) \geqslant 0$,

$$\underbrace{\left(\frac{\gamma+h}{2}\frac{\partial F(s_H^*(\lambda);\sigma)}{\partial s_H^*(\lambda)} + \frac{t}{\lambda}\right)}_{>0}\frac{ds_H^*(\lambda)}{d\lambda} = \underbrace{\frac{t}{\lambda^2}(\Delta + 1/2)}_{>0} - \underbrace{\frac{\gamma+h}{2}\frac{\partial F(s_H^*(\lambda);\sigma)}{\partial \sigma}\frac{d\sigma}{d\lambda}}_{\leqslant 0}.$$

Thus we again have $\frac{ds_H^*(\lambda)}{d\lambda} > 0$, and since by Lemma 10 $s_H^*(\lambda) = -s(\mathbf{p}^*(\lambda), \lambda) < 0$, we have shown that $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ for $i \in \{H, L\}$.

(1) and (2) together lead to the first result of the proposition, i.e., $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ for $i \in \{H, L\}$ for any $\lambda > \hat{\lambda}$.

We now prove the second result of the proposition, $\frac{d\bar{L}(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$. From Equation (L), $\bar{L}(\mathbf{p}, \lambda) = (\gamma + h)\left[\int_{-\infty}^{-s(\mathbf{p},\lambda)} F(\epsilon; \sigma) d\epsilon + \frac{s(\mathbf{p},\lambda)}{2}\right]$, and thus

$$\frac{d\bar{L}(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} = \frac{\partial \bar{L}(\mathbf{p}^*(\lambda),\lambda)}{\partial s(\mathbf{p}^*(\lambda),\lambda)}\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} + \frac{d\sigma}{d\lambda} \cdot \int_{-\infty}^{-s(\mathbf{p}^*(\lambda),\lambda)} \frac{F(\epsilon;\sigma)}{\sigma} d\epsilon.$$

The first term is strictly negative, since $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$ and $\frac{\partial \bar{L}(\mathbf{p},\lambda)}{\partial s(\mathbf{p},\lambda)} = (\gamma + h)\left[\frac{1}{2} - F(-s(\mathbf{p}, \lambda); \sigma)\right] > 0$. We used the fact that $s(\mathbf{p}, \lambda) > 0$ by Lemma 10. The second term is non-positive, since $\frac{F(\epsilon;\sigma)}{\sigma} = -\frac{\epsilon}{\sigma} f(\epsilon; \sigma) \geqslant 0$ for all $\epsilon \leqslant 0$, and $\frac{d\sigma}{d\lambda} \leqslant 0$. Thus, $\frac{d\bar{L}(\mathbf{p}^*(\lambda),\lambda)}{d\lambda} < 0$.

D.4. Proof of Proposition 3

Recall from Equation (L) that the expected loss averaged over time is

$$\bar{L}(\mathbf{p}, \lambda) = (\gamma + h)\left[\int_{-\infty}^{-s(\mathbf{p},\lambda)} F(\epsilon; \sigma) d\epsilon + \frac{s(\mathbf{p},\lambda)}{2}\right] \equiv (\gamma + h)\left[\int_{-\infty}^{s(\mathbf{p},\lambda)} F(\epsilon; \sigma) d\epsilon - \frac{s(\mathbf{p},\lambda)}{2}\right],$$

where $s(\mathbf{p}, \lambda) = |\Delta - \frac{\lambda}{t} p_\Delta(\lambda)| = \Delta - \frac{\lambda}{t} p_\Delta(\lambda)$ by Lemma 10. Thus, reduction in the expected loss is given by

$$\bar{L}(\bar{\mathbf{p}}, 0) - \bar{L}(\mathbf{p}^*(\lambda), \lambda) = (\gamma + h)\left[\int_{s(\mathbf{p}^*(\lambda),\lambda)}^{\Delta} F(\epsilon; \sigma) d\epsilon - \frac{1}{2}[\Delta - s(\mathbf{p}^*(\lambda), \lambda)]\right].$$

We first show that $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \Delta} \geqslant 0$. Differentiating,

$$\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \Delta} = (\gamma + h)\left[F(\Delta; \sigma) - \frac{1}{2} - \frac{\partial s(\mathbf{p}^*(\lambda),\lambda)}{\partial \Delta}\left(F(s(\mathbf{p}^*(\lambda), \lambda); \sigma) - \frac{1}{2}\right)\right].$$

- When $p_H^*(\lambda) > \underline{p}$ and $p_L^*(\lambda) = \underline{p}$, Lemma 10 and Equation (FOC-2) that $s_H^*(\lambda) < 0$ and satisfies the first order condition, $\frac{t}{\lambda} s_H^*(\lambda) = -\frac{\gamma+h}{2} F(s_H^*(\lambda); \sigma) + \frac{1}{2}(1 - \underline{p} + \gamma) - \frac{t}{\lambda}(1/2 + \Delta)$. Differentiating both sides with respect to Δ and rearranging,

$$\frac{\partial s(\mathbf{p}^*(\lambda),\lambda)}{\partial \Delta} = -\frac{\partial \partial_H^*(\lambda)}{\partial \Delta} = \frac{t}{\lambda} \cdot \left(\frac{t}{\lambda} + \frac{\gamma+h}{2} f(s_H^*(\lambda); \sigma)\right)^{-1} \leqslant 1.$$

Since $F(s(\mathbf{p}^*(\lambda), \lambda); \sigma) - \frac{1}{2} > 0$ by Lemma 10, $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \Delta} \geqslant (\gamma + h)\left[F(\Delta; \sigma) - F(s(\mathbf{p}^*(\lambda), \lambda); \sigma)\right] \geqslant 0$. If $f(s_H^*; \sigma) = f(\Delta - \frac{\lambda}{t} p_\Delta^*(\lambda); \sigma) > 0$, $\frac{\partial s(\mathbf{p}^*(\lambda),\lambda)}{\partial \Delta} < 1$ and the inequality is strict.

- When $p_L^*(\lambda) > \underline{p}$, using the expression for p_Δ^* in Equation (7), and noting that $s_H^*(\lambda) = \frac{\lambda}{t}(p_H^*(\lambda) - p_L^*(\lambda)) - \Delta$, we have $\frac{t}{\lambda}s_H^*(\lambda) = \frac{\gamma+h}{3}[1-2F(s_H^*(\lambda);\sigma)] - \frac{t}{\lambda}\Delta$. Differentiating both sides with respect to Δ and rearranging,

$$\frac{\partial s(\mathbf{p}^*(\lambda),\lambda)}{\partial \Delta} = -\frac{\partial s_H^*(\lambda)}{\partial \Delta} = \frac{t}{\lambda}\cdot\left(\frac{t}{\lambda} + \frac{2(\gamma+h)}{3}f(s_H^*(\lambda);\sigma)\right)^{-1} \leqslant 1.$$

  Similarly, $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \Delta} \geqslant (\gamma+h)\left[F(\Delta;\sigma) - F(s(\mathbf{p}^*(\lambda),\lambda);\sigma)\right] \geqslant 0$. If $f(s_H^*;\sigma) = f(\Delta - \frac{\lambda}{t}p_\Delta^*(\lambda);\sigma) > 0$, $\frac{\partial s(\mathbf{p}^*(\lambda),\lambda)}{\partial \Delta} < 1$ and the inequality is strict.

Next, we show that $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \sigma} \leqslant 0$. Differentiating,

$$\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \sigma} = (\gamma+h)\left[-\frac{1}{\sigma}\int_{s(\mathbf{p}^*(\lambda),\lambda)}^{\Delta} \epsilon f(\epsilon;\sigma)d\epsilon - \frac{\partial s(\mathbf{p}^*(\lambda),\lambda)}{\partial \sigma}\left(F(s(\mathbf{p}^*(\lambda),\lambda);\sigma) - \frac{1}{2}\right)\right].$$

Using the second result from Proposition 1, $\frac{ds(\mathbf{p}^*(\lambda),\lambda)}{d\sigma} = -\frac{\lambda}{t}\frac{\partial p_\Delta^*(\lambda)}{\partial \sigma} \geqslant 0$. Thus, it is easy to verify that $\frac{\partial \bar{L}_\Delta(\lambda)}{\partial \sigma} \leqslant 0$ with strict inequality when $f(\Delta - \frac{\lambda}{t}p_\Delta^*(\lambda);\sigma) > 0$.

D.5. Proof of Theorem 2

By Lemma 1, all sellers strictly benefit if $\lambda > \hat{\lambda}$. We thus focus on showing that buyers strictly benefit if $\Delta > \underline{\Delta}$, where $F(\underline{\Delta};\sigma) = \frac{\gamma+1-\underline{p}}{\gamma+h}$. The condition, $h \geqslant 1-\underline{p}$, is needed such that $\underline{\Delta}$ is defined.

We make use of Lemma 11 below, which establishes that the low-demand buyer is not worse off at $p_H > \underline{p}$ compared to $p_H = \underline{p}$, if $p_H \leqslant \hat{p}$, where $\hat{p}$ satisfies $F\left(\Delta - \frac{\lambda}{t}(\hat{p} - \underline{p});\sigma\right) \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h}$.

Lemma 11. *For any $\lambda > 0$ and any $\hat{p} > \underline{p}$, if $F\left(\Delta - \frac{\lambda}{t}(\hat{p} - \underline{p});\sigma\right) \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h}$, then $\pi_L(\underline{p}, p_H, \lambda) \geqslant \pi_L(\underline{p}, \underline{p}, \lambda)$ for all $p_H \in [\underline{p}, \hat{p}]$.*

Proof. Differentiating $\pi_L(p_L, p_H, \lambda)$ with respect to p_H,

$$\left.\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\right|_{p_L=\underline{p}} = -\frac{\lambda}{t}\left[1-\underline{p} + \gamma[1 - F(\Delta - \frac{\lambda}{t}(p_H - \underline{p});\sigma)] - hF(\Delta - \frac{\lambda}{t}(p_H - \underline{p};\sigma)\right].$$

Thus, $\left.\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\right|_{p_L=\underline{p}} \geqslant 0 \iff F\left(\Delta - \frac{\lambda}{t}(p_H - \underline{p});\sigma\right) \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h}$. Now, since at $p_H = \hat{p}$, we have $F\left(\Delta - \frac{\lambda}{t}(\hat{p} - \underline{p});\sigma\right) \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h}$, this implies $F\left(\Delta - \frac{\lambda}{t}(p_H - \underline{p});\sigma\right) \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h}$ at all $p_H \in [\underline{p}, \hat{p}]$ since $F\left(\Delta - \frac{\lambda}{t}(p_H - \underline{p});\sigma\right)$ increases with decreasing p_H. Consequently, $\left.\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\right|_{p_L=\underline{p}} \geqslant 0$ at all $p_H \in [\underline{p}, \hat{p}]$. Hence, $\pi_L(\underline{p}, p_H, \lambda) = \pi_L(\underline{p}, \underline{p}, \lambda) + \int_{\underline{p}}^{p_H} \left.\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\right|_{p_L=\underline{p}} dp_H \geqslant \pi_L(\underline{p}, \underline{p}, \lambda)$ for all $p_H \in [\underline{p}, \hat{p}]$. □

For any $\Delta > \underline{\Delta}$, let $\hat{p} = \underline{p} + \frac{t}{\lambda}(\Delta - \underline{\Delta}) > \underline{p}$. Then, for any λ such that the competitive equilibrium satisfies $p_H^*(\lambda) \in (\underline{p}, \hat{p}]$, $p_L^*(\lambda) = \underline{p}$, we have

1. the high-demand buyer is strictly better off, i.e., $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) > \pi_H(\underline{p}, \underline{p}, \lambda) = \pi_H(\underline{p}, \underline{p}, 0)$. This is because $p_H^*(\lambda) = p_H^{FOC}(\underline{p}, \lambda)$ (see Case 3 of Appendix C.1), and $\pi_H(p_H, p_L, \lambda)$ is strictly concave in p_H.
2. The low-demand buyer is weakly better off, i.e., $\pi_L(\underline{p}, p_H^*(\lambda), \lambda) \geqslant \pi_L(\underline{p}, \underline{p}, \lambda) = \pi_L(\underline{p}, \underline{p}, 0)$ (Lemma 11).

Thus, buyers are strictly better off in expectation, i.e., $\bar{\pi}(\mathbf{p}^*(\lambda), \lambda) > \bar{\pi}(\underline{\mathbf{p}}, 0)$.

Given the full characterization of the equilibrium in Appendix C.1, we have the following observation. For any $\Delta > \underline{\Delta}$, let $\hat{p} = \underline{p} + \frac{t}{\lambda}(\Delta - \underline{\Delta}) > \underline{p}$. Then, there always exists $\lambda' > \hat{\lambda}$, such that for any $\lambda \in (\hat{\lambda}, \lambda')$, the competitive equilibrium prices satisfy $p_H^*(\lambda) \in (\underline{p}, \hat{p})$ and $p_L^*(\lambda) = \underline{p}$, where $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma) - h + 1 - \underline{p}}$. Consequently, strong Pareto improvement arises for $\lambda \in (\hat{\lambda}, \lambda')$, which leads to the result in Theorem 2.

Extension to $\frac{d\sigma}{d\lambda} \leqslant 0$. Theorem 2 assumes that supply uncertainty is a constant σ independent of λ, i.e., $\frac{d\sigma}{d\lambda} = 0$. It is straightforward to verify that the result in Theorem 2 still holds when $\frac{d\sigma}{d\lambda} \leqslant 0$, i.e., we allow σ to be either independent or non-increasing in λ, with $\underline{\Delta}$ defined by the implicit function $F(\underline{\Delta};\sigma(0)) = \frac{\gamma+1-p}{\gamma+h}$. The proof is similar, and the only main difference is that instead of $\pi_i(\underline{p},\underline{p},\lambda) = \pi_i(\underline{p},\underline{p},0)$, we now have the inequality $\pi_i(\underline{p},\underline{p},\lambda) \geqslant \pi_i(\underline{p},\underline{p},0)$ since decreasing σ reduces the expected cost incurred by buyers.

D.6. Proof of Proposition 4

It is easy to verify that $\hat{\lambda} = \frac{t}{(\gamma+h)F(\Delta;\sigma)-h+1-\underline{p}}$ is (weakly) decreasing in Δ, γ and increasing in σ, t. It remains to show that if $p_L^*(\lambda^*) > \underline{p}$, then λ^* is decreasing in σ, γ and increasing in Δ, t. We show this by showing that the change in buyers' expected profit, $\bar{\pi}(\mathbf{p}^*(\lambda),\lambda) - \bar{\pi}(\underline{\mathbf{p}},0)$, is decreasing in σ, γ and increasing in Δ, t for all λ such that $p_L^*(\lambda) > \underline{p}$. Denoting $\pi_\Delta(\lambda) = \bar{\pi}(\mathbf{p}^*(\lambda),\lambda) - \bar{\pi}(\underline{\mathbf{p}},0)$, we will show that for all λ such that $p_L^*(\lambda) > \underline{p}$, $\frac{\partial \pi_\Delta(\lambda)}{\partial \sigma} \leqslant 0$, $\frac{\partial \pi_\Delta(\lambda)}{\partial \gamma} \leqslant 0$, $\frac{\partial \pi_\Delta(\lambda)}{\partial \Delta} \geqslant 0$, $\frac{\partial \pi_\Delta(\lambda)}{\partial t} \geqslant 0$. The results in the proposition follows.

Using profit Equations (RP), (L) and equilibrium prices when $p_L^*(\lambda) > \underline{p}$ described by Equations (6) and (7) in Theorem 1, the change in buyer profit can be written as

$$\pi_\Delta(\lambda) = \frac{h-\gamma}{2} + \frac{t}{\lambda} - \lambda\frac{p_\Delta^{*2}(\lambda)}{2t} - (1-\underline{p}) + (\gamma+h)\left[\int_{\Delta-\frac{\lambda p_\Delta^*(\lambda)}{t}}^{\Delta} F(\epsilon;\sigma)d\epsilon - \frac{\lambda p_\Delta^*(\lambda)}{2t}\right],$$

where $p_\Delta^*(\lambda)$ satisfies the equality $(\gamma+h)\left[F\left(\Delta - \frac{\lambda p_\Delta^*(\lambda)}{t};\sigma\right) - \frac{1}{2}\right] = \frac{3}{2}p_\Delta^*(\lambda)$. Observe that, using the equality for $p_\Delta^*(\lambda)$, we have $\frac{\partial \pi_\Delta}{\partial p_\Delta^*} = -\frac{\lambda}{t}\left(p_\Delta^* - (\gamma+h)\left[F\left(\Delta - \frac{\lambda p_\Delta^*}{t};\sigma\right) - \frac{1}{2}\right]\right) = \frac{\lambda}{2t}p_\Delta^* \geqslant 0$.

1. We show $\frac{\partial \pi_\Delta}{\partial \sigma} \leqslant 0$. Differentiating $\pi_\Delta(\lambda)$ with respect to σ,

$$\frac{\partial \pi_\Delta}{\partial \sigma} = \frac{\partial \pi_\Delta}{\partial p_\Delta^*}\frac{\partial p_\Delta^*}{\partial \sigma} - \frac{\gamma+h}{\sigma}\int_{\Delta-\frac{\lambda p_\Delta^*(\lambda)}{t}}^{\Delta} \epsilon f(\epsilon;\sigma)d\epsilon \leqslant \frac{\partial \pi_\Delta}{\partial p_\Delta^*}\frac{\partial p_\Delta^*}{\partial \sigma} = \frac{\lambda}{t}\frac{p_\Delta^*}{2}\frac{\partial p_\Delta^*}{\partial \sigma} \leqslant 0.$$

The first $\leqslant$ follows from the fact that $\Delta - \frac{\lambda p_\Delta^*(\lambda)}{t} \geqslant 0$ (Lemma 10) and thus $-\frac{\gamma+h}{\sigma}\int_{\Delta-\frac{\lambda p_\Delta^*}{t}}^{\Delta} \epsilon f(\epsilon;\sigma)d\epsilon \leqslant 0$. The second equality uses $\frac{\partial \pi_\Delta}{\partial p_\Delta^*} = \frac{\lambda}{2t}p_\Delta^* \geqslant 0$. The last $\leqslant$ follows from the fact that $\frac{\partial p_\Delta^*}{\partial \sigma} \leqslant 0$ (Proposition 1).

2. We show $\frac{\partial \pi_\Delta}{\partial t} \geqslant 0$. We first show that $\frac{\partial p_\Delta^*}{\partial t} \geqslant 0$. Differentiating Equation (7), we have

$$(\gamma+h)\lambda f(\Delta - \lambda p_\Delta^*/t;\sigma)\left(-\frac{1}{t}\frac{\partial p_\Delta^*}{\partial t} + \frac{p_\Delta^*}{t^2}\right) = \frac{3}{2}\frac{\partial p_\Delta^*}{\partial t} \quad \Rightarrow \quad \frac{\partial p_\Delta^*}{\partial t} = \frac{\lambda f(\Delta - \lambda p_\Delta^*/t;\sigma)p_\Delta^*/t^2}{\frac{\lambda f(\Delta-\lambda p_\Delta^*/t;\sigma)}{t} + \frac{3}{2(\gamma+h)}} \in \left[0, \frac{p_\Delta^*}{t}\right].$$

Differentiating $\pi_\Delta(\lambda)$ with respect to t,

$$\frac{\partial \pi_\Delta}{\partial t} = \frac{1}{\lambda} + \lambda\frac{p_\Delta^{*2}}{2t^2} - \frac{\lambda p_\Delta^*}{t}\frac{\partial p_\Delta^*}{\partial t} + \lambda(\gamma+h)\left[F\left(\Delta - \frac{\lambda p_\Delta^*}{t};\sigma\right) - \frac{1}{2}\right]\left(\frac{1}{t}\frac{\partial p_\Delta^*}{\partial t} - \frac{p_\Delta^*}{t^2}\right) = \frac{1}{\lambda} - \lambda\frac{p_\Delta^{*2}}{t^2} + \frac{\lambda p_\Delta^*}{2t}\frac{\partial p_\Delta^*}{\partial t} \geqslant 0.$$

The second equality follows from substituting $(\gamma+h)\left[F\left(\Delta - \frac{\lambda p_\Delta^*(\lambda)}{t};\sigma\right) - \frac{1}{2}\right] = \frac{3}{2}p_\Delta^*(\lambda)$ and simplifying. The last inequality holds since $\lambda \leqslant 1$, $p_\Delta^* \leqslant t$, and $\frac{\partial p_\Delta^*}{\partial t} \geqslant 0$.

(3) We show $\frac{\partial \pi_\Delta}{\partial \gamma} \leqslant 0$ if $\Delta \leqslant \frac{3}{4}$. Differentiating Equation (7) and rearranging,

$$\frac{\partial p_\Delta^*}{\partial \gamma} = \frac{F\left(\Delta - \frac{\lambda p_\Delta^*}{t};\sigma\right) - \frac{1}{2}}{(\gamma+h)\frac{\lambda}{t}f(\Delta - \lambda p_\Delta^*/t;\sigma) + \frac{3}{2}} \leqslant \frac{1}{3}\left[2F\left(\Delta - \frac{\lambda p_\Delta^*}{t};\sigma\right) - 1\right] \leqslant \frac{1}{3}.$$

Differentiating $\pi_\Delta(\lambda)$ with respect to γ,

$$\frac{\partial \pi_\Delta}{\partial \gamma} = \frac{\partial \pi_\Delta}{\partial p_\Delta^*}\frac{\partial p_\Delta^*}{\partial \gamma} - \frac{1}{2} + \int_{\Delta-\frac{\lambda p_\Delta^*}{t}}^{\Delta} F(\epsilon;\sigma)d\epsilon - \frac{\lambda p_\Delta^*}{2t} \leqslant \frac{2\lambda}{3t}p_\Delta^* - \frac{1}{2} \leqslant 0.$$

The first $\leqslant$ uses the fact that $\frac{\partial \pi_\Delta}{\partial p^*_\Delta} = \frac{\lambda}{2t} p^*_\Delta$, $\frac{\partial p^*_\Delta}{\partial \gamma} \leqslant \frac{1}{3}$, and $\int_{\Delta - \frac{\lambda p^*_\Delta}{t}}^{\Delta} F(\epsilon;\sigma) d\epsilon \leqslant \frac{\lambda}{t} p^*_\Delta$. The second $\leqslant$ uses the fact that $\Delta \leqslant \frac{3}{4}$, and $\frac{\lambda}{t} p^*_\Delta < \Delta \leqslant \frac{3}{4}$ (Lemma 10).

(4) We show $\frac{\partial \pi_\Delta}{\partial \Delta} \geqslant 0$. Differentiating Equation (7) and rearranging,

$$(\gamma+h) f(\Delta - \lambda p^*_\Delta / t; \sigma) \left(1 - \frac{\lambda}{t} \frac{\partial p^*_\Delta}{\partial \Delta}\right) = \frac{3}{2} \frac{\partial p^*_\Delta}{\partial \Delta} \quad \Rightarrow \quad \frac{\partial p^*_\Delta}{\partial \Delta} = \frac{f(\Delta - \lambda p^*_\Delta/t;\sigma)}{\lambda \frac{f(\Delta - \lambda p^*_\Delta/t;\sigma)}{t} + \frac{3}{2(\gamma+h)}} \geqslant 0.$$

Differentiating $\pi_\Delta(\lambda)$ with respect to Δ, $\frac{\partial \pi_\Delta}{\partial \Delta} = \frac{\partial \pi_\Delta}{\partial p^*_\Delta} \frac{\partial p^*_\Delta}{\partial \Delta} + (\gamma+h)[F(\Delta;\sigma) - F(\Delta - \lambda p^*_\Delta/t;\sigma)] \geqslant 0$.

D.7. Proof of Proposition 5

The joint profit for given market prices $\bar{p}$, p_Δ under transparency λ is given by Equation (P-joint) in Appendix B.1.1, $\pi_H(\bar{p}, p_\Delta, \lambda) + \pi_L(\bar{p}, p_\Delta, \lambda) = 2 - 2\bar{p} - \frac{\lambda}{t} p_\Delta^2 - (\gamma+h)\left[2\int_{-\infty}^{\Delta - \frac{\lambda}{t} p_\Delta} F(\epsilon;\sigma) d\epsilon - \Delta + \frac{\lambda}{t} p_\Delta\right]$.

Let $p_H^{c*}(\lambda), p_L^{c*}(\lambda)$ denote the jointly profit-maximizing collusive prices, and let $\bar{p}^{c*}(\lambda), p_\Delta^{c*}(\lambda)$ denote the average price and price difference given these prices. It is straightforward to check using the above equation that the jointly profit-maximizing collusive prices $p_H^{c*}(\lambda), p_L^{c*}(\lambda)$ must satisfy $p_L^{c*}(\lambda) = \underline{p}$ for any λ, i.e., $\bar{p}^{c*}(\lambda) = \underline{p} + p_\Delta^{c*}(\lambda)/2$, and $p_\Delta^{c*}(\lambda) \geqslant 0$. The joint profit-maximizing problem is thus

$$p_\Delta^{c*}(\lambda) = \underset{p_\Delta \geqslant 0}{\arg\max}\ g(p_\Delta, \lambda), \quad \text{where } g(p_\Delta, \lambda) = 2\bar{\pi}(\mathbf{p}^c, \lambda)\,|_{p_L^c = \underline{p}} \text{ satisfies}$$

$$g(p_\Delta, \lambda) = 2(1 - \underline{p} - p_\Delta/2) - \frac{\lambda}{t} p_\Delta^2 - (\gamma+h)\left[2\int_{-\infty}^{\Delta - \frac{\lambda}{t} p_\Delta} F(\epsilon;\sigma) d\epsilon - \Delta + \frac{\lambda}{t} p_\Delta\right].$$

Since $g(p_\Delta, \lambda)$ is continuously differentiable and strictly concave in p_Δ, the optimal solution $p_\Delta^{c*}(\lambda)$ only needs to satisfy the first order condition, $\frac{\partial g(p_\Delta,\lambda)}{\partial p_\Delta} = 0$. Differentiating,

$$\frac{\partial g(p_\Delta, \lambda)}{\partial p_\Delta} = -1 - \frac{2\lambda}{t} p_\Delta + \frac{\lambda}{t}(\gamma+h)\left[2F\left(\Delta - \frac{\lambda}{t} p_\Delta; \sigma\right) - 1\right], \quad \text{and}$$

$$\left.\frac{\partial g(p_\Delta, \lambda)}{\partial p_\Delta}\right|_{p_\Delta = 0} = -1 + \frac{\lambda}{t}(\gamma+h)[2F(\Delta;\sigma) - 1].$$

We have

$$p_\Delta^{c*}(\lambda) \begin{cases} = 0 & \text{if } \frac{\lambda}{t}(\gamma+h)[2F(\Delta;\sigma) - 1] \leqslant 1, \\ > 0 & \text{otherwise.} \end{cases}$$

When $p_\Delta^{c*}(\lambda) > 0$, it satisfies the first order condition, $p_\Delta^{c*} = (\gamma+h)\left[F(\Delta - \frac{\lambda}{t} p_\Delta^{c*};\sigma) - \frac{1}{2}\right] - \frac{t}{2\lambda}$. Note that this equation always has a unique solution since the RHS is monotonically decreasing in p_Δ^{c*}. The result in Proposition 5 follows. Note that, as explained at the end of Section 3, we only focus on the case where $p_\Delta^{c*}(\lambda)$ obtained from the first-order condition above does not exceed t. If $p_\Delta^{c*}(\lambda) > t$, the optimal price dispersion will be capped at t, i.e., $p_L^{c*} = \underline{p}$ and $p_H^{c*} = \underline{p} + t$.

D.8. Proof of Proposition 6

By Equation (12) in Proposition 5, the average market price $\bar{p}^{c*}$ is higher than the reservation price $\underline{p}$ if and only if $\lambda > \hat{\lambda}^c = \frac{t}{(\gamma+h)[2F(\Delta;\sigma)-1]}$. Thus, all sellers strictly benefit if and only if $\lambda > \hat{\lambda}^c$. For buyers, optimality and uniqueness of $\mathbf{p}^{c*}$ (Proposition 5) implies $\bar{\pi}(\mathbf{p}^{c*}(\lambda), \lambda) > \bar{\pi}(\underline{\mathbf{p}}, \lambda) = \bar{\pi}(\underline{\mathbf{p}}, 0)$ if and only if $\lambda > \hat{\lambda}^c$. Now since $\lambda \in [0,1]$, we need $\hat{\lambda}^c = \frac{t}{(\gamma+h)[2F(\Delta;\sigma)-1]} < 1$, or, equivalently, $\gamma + h > t$ and $F(\Delta;\sigma) > \frac{1}{2} + \frac{t}{2(\gamma+h)}$. This gives the result in Proposition 6.

Extension to $\frac{d\sigma}{d\lambda} \leqslant 0$. Proposition 6 assumes that supply uncertainty is independent of λ, $\frac{d\sigma}{d\lambda} = 0$. We now relax the assumption of independence to $\frac{d\sigma}{d\lambda} \leqslant 0$, i.e., we allow σ to be either independent or non-increasing in λ. It is straightforward to verify that the result in Proposition 6 still holds, with $\underline{\Delta}^c$ defined by the implicit function $F(\underline{\Delta}^c; \sigma(0)) = \frac{t+\gamma+h}{2(\gamma+h)}$ and $\hat{\lambda}^c$ defined by the implicit function $\frac{\hat{\lambda}^c}{t}(\gamma+h)[2F(\Delta;\sigma(\hat{\lambda}^c)) - 1] = 1$. The proof is similar, and the only main difference is that instead of $\bar{\pi}(\mathbf{p}^{c*}, \lambda) > \bar{\pi}(\underline{\mathbf{p}}, \lambda) = \bar{\pi}(\underline{\mathbf{p}}, 0)$, we now have the inequality $\bar{\pi}(\mathbf{p}^{c*}, \lambda) > \bar{\pi}(\underline{\mathbf{p}}, \lambda) \geqslant \bar{\pi}(\underline{\mathbf{p}}, 0)$ due to reduced supply uncertainty.

D.9. Proof of Theorem 3

Theorem 3 assumes that supply uncertainty is independent of λ, $\frac{d\sigma}{d\lambda} = 0$. Here, we present the proof for a more general version of the theorem, where we relax the assumption of independence to $\frac{d\sigma}{d\lambda} \leqslant 0$, i.e., we allow σ to be either independent or non-increasing in λ. We show that the result in Theorem 3 holds as long as $\frac{d\sigma}{d\lambda} \leqslant 0$ is satisfied, and the original proposition is a special case where $\frac{d\sigma}{d\lambda} = 0$. We will prove this result by proving the three lemmas below, under the assumption that $\underline{\mathbf{p}}$ is a collusive equilibrium.

Lemma 12 characterizes a continuum of collusive outcomes that are more profitable than $\underline{\mathbf{p}}$ and are strongly Pareto-improving for $\lambda > \hat{\lambda}^c$.

Lemma 12. *If $\Delta > 0$ and $\frac{d\sigma}{d\lambda} \leqslant 0$ for any λ, any collusive outcome $\mathbf{p}^c = (p_H^c, p_L^c)$ with $\underline{p} < p_H^c \leqslant p_H^{c*}$ and $p_L^c = \underline{p}$ is strictly more profitable in expectation than $\underline{\mathbf{p}}$ and strongly Pareto-improving for any $\lambda > \hat{\lambda}^c$.*

Proof. First, the sellers are always strictly better off in expectation if $p_H^c > \underline{p}$. It remains to show that the buyers are always strictly better off both compared to collusion at the reservation price and compared to the original market with no price transparency. Specifically, we show that, for any $\lambda > \hat{\lambda}^c$, for any collusive outcome $\mathbf{p}^c$ where $\underline{p} < p_H^c \leqslant p_H^{c*}$ and $p_L^c = \underline{p}$, it holds that $\bar{\pi}(\mathbf{p}^c, \lambda) > \bar{\pi}(\underline{\mathbf{p}}, \lambda) \geqslant \bar{\pi}(\underline{\mathbf{p}}, 0)$.

Observe that for any $\mathbf{p}^c$ such that $p_L^c = \underline{p}$, the expected buyer profit $\bar{\pi}(\mathbf{p}^c, \lambda)$ satisfies

$$\frac{\partial \bar{\pi}(\mathbf{p}^c, \lambda)}{\partial p_\Delta} \begin{cases} > 0 & \text{if } 0 < p_\Delta < p_\Delta^{c*}(\lambda) \\ = 0 & \text{if } p_\Delta = p_\Delta^{c*}(\lambda), \end{cases}$$

for all $\lambda > \hat{\lambda}^c$. Therefore, $\bar{\pi}(\mathbf{p}^c, \lambda) > \bar{\pi}(\underline{\mathbf{p}}, \lambda)$ for any $\mathbf{p}^c$ where $\underline{p} < p_H^c \leqslant p_H^{c*}$ and $p_L^c = \underline{p}$.

We now compare $\bar{\pi}(\underline{\mathbf{p}}, \lambda)$ and $\bar{\pi}(\underline{\mathbf{p}}, 0)$. Under the cooperative strategy $\underline{\mathbf{p}}$, if $\frac{d\sigma}{d\lambda} \leqslant 0$ then $DS = 0$ and $SV(\lambda) \geqslant 0$, and so $\bar{L}(\underline{\mathbf{p}}, \lambda) \leqslant L(\underline{\mathbf{p}}, 0)$. In addition, $\bar{\Psi}(\underline{\mathbf{p}}, \lambda) = \bar{\Psi}(\underline{\mathbf{p}}, 0) = 1 - \underline{\mathbf{p}}$. Thus

$$\bar{\pi}(\underline{\mathbf{p}}, \lambda) = \bar{\Psi}(\underline{\mathbf{p}}, \lambda) - \bar{L}(\underline{\mathbf{p}}, \lambda) \geqslant \bar{\Psi}(\underline{\mathbf{p}}, 0) - \bar{L}(\underline{\mathbf{p}}, 0) \geqslant \bar{\pi}(\underline{\mathbf{p}}, 0),$$

as required. □

The next two lemmas prove existence and also present results regarding stability of collusion. We use the following formal definition and notation for the buyers' threshold discount factors.

Definition 1. *For $i \in \{H, L\}$, prices $\mathbf{p}^c$, and transparency level λ, we define $\pi_i^d(\mathbf{p}^c, \lambda) = \max_{p_i} \pi_i(p_i, p_j^c, \lambda)$ as the maximum single-period expected profit in the case of a single-period defect by buyer i, and let*

$$\hat{\delta}_i(\mathbf{p}^c, \lambda) = \frac{\pi_i^d(\mathbf{p}^c, \lambda) - \pi_i(\mathbf{p}^c, \lambda)}{\pi_i^d(\mathbf{p}^c, \lambda) - \pi_i(\mathbf{p}^c, \lambda) + \bar{\pi}(\mathbf{p}^c, \lambda) - \bar{\pi}(\mathbf{p}^*(\lambda), \lambda)} \tag{22}$$

be the threshold discount factor for buyer i that satisfies the non-deviation constraint (13)*. The threshold discount factor for the collusive outcome $\mathbf{p}^c$ is thus $\hat{\delta}(\mathbf{p}^c, \lambda) = \max\{\hat{\delta}_H(\mathbf{p}^c, \lambda), \hat{\delta}_L(\mathbf{p}^c, \lambda)\}$.*

Lemma 13 states that under the collusive outcome $\underline{\mathbf{p}}$, buyer H always has greater incentive to defect from collusion compared to buyer L, and therefore the threshold discount factor of buyer H determines the overall threshold discount factor, i.e., $\hat{\delta}(\mathbf{p}^c,\lambda)=\hat{\delta}_H(\mathbf{p}^c,\lambda)$.

Lemma 13. *If* $\Delta>0$, $\hat{\delta}_L(\underline{\mathbf{p}},\lambda)<\hat{\delta}_H(\underline{\mathbf{p}},\lambda)$ *for all* $\lambda>\hat{\lambda}$, *and hence* $\hat{\delta}(\underline{\mathbf{p}},\lambda)=\hat{\delta}_H(\underline{\mathbf{p}},\lambda)$ *for all* $\lambda>\hat{\lambda}$.

Proof. First, it is obvious that buyer H always has the incentive to deviate from $p_H=\underline{p}$ for any $\lambda>\hat{\lambda}$ since $p_H^{FOC}(\underline{p},\lambda)>\underline{p}$ for all $\hat{\lambda}<\lambda<1$. When buyer L has no incentive to deviate and H does, $\hat{\delta}_H>\hat{\delta}_L=0$. We now focus on the case where both H and L have the incentive to deviate, i.e., $\hat{\delta}_L>0$ and $\hat{\delta}_H>0$.

Under $\underline{\mathbf{p}}$, using Definition 22, we see that to show that $\hat{\delta}_L(\underline{\mathbf{p}},\lambda)<\hat{\delta}_H(\underline{\mathbf{p}},\lambda)$, it suffices to show

$$\pi_L^d(\underline{\mathbf{p}},\lambda)-\pi_L(\underline{\mathbf{p}},\lambda)<\pi_H^d(\underline{\mathbf{p}},\lambda)-\pi_H(\underline{\mathbf{p}},\lambda). \tag{23}$$

We now show that Equation (23) is true. When both buyers collude at $\underline{p}$, the buyer's optimal deviation strategies in H,L demand states are given by $p_H^{FOC}(\underline{p},\lambda),p_L^{FOC}(\underline{p},\lambda)>\underline{p}$ respectively, i.e.,

$$\pi_H^d(\underline{\mathbf{p}},\lambda)=\pi_H(p_H^{FOC}(\underline{p},\lambda),\underline{p},\lambda),\quad \pi_L^d(\underline{\mathbf{p}},\lambda)=\pi_L(p_L^{FOC}(\underline{p},\lambda),\underline{p},\lambda).$$

In addition, let $\pi_H'^d(\underline{\mathbf{p}},\lambda)$ denote buyer H's deviation profit when pricing at $p_L^{FOC}(\underline{p},\lambda)$ instead of $p_H^{FOC}(\underline{p},\lambda)$:

$$\pi_H'^d(\underline{\mathbf{p}},\lambda)=\pi_H(p_L^{FOC}(\underline{p},\lambda),\underline{p},\lambda)\leqslant\pi_H^d(\underline{\mathbf{p}},\lambda),$$

where the inequality holds because of the optimality of π_H^d. Thus we have

$$[\pi_L^d(\underline{\mathbf{p}},\lambda)-\pi_L(\underline{\mathbf{p}},\lambda)]-[\pi_H^d(\underline{\mathbf{p}},\lambda)-\pi_H(\underline{\mathbf{p}},\lambda)]\leqslant[\pi_L^d(\underline{\mathbf{p}},\lambda)-\pi_L(\underline{\mathbf{p}},\lambda)]-[\pi_H'^d(\underline{\mathbf{p}},\lambda)-\pi_H(\underline{\mathbf{p}},\lambda)].$$

Let $p_L^{FOC}(\underline{p},\lambda)=\underline{p}+p_\Delta(\lambda)$, $\bar{p}=\frac{p_L^{FOC}(\underline{p},\lambda)+\underline{p}}{2}=\underline{p}+\frac{p_\Delta(\lambda)}{2}$, and let $F(\cdot)$ be a shorthand for the CDF, $F(\epsilon;\sigma)$. Expanding the RHS and simplifying using the expressions for π_H,π_L in Equations (P-H)-(P-joint) yields

$$\begin{aligned}
&[\pi_L^d(\underline{\mathbf{p}},\lambda)-\pi_L(\underline{\mathbf{p}},\lambda)]-[\pi_H'^d(\underline{\mathbf{p}},\lambda)-\pi_H(\underline{\mathbf{p}},\lambda)]\\
=&[\pi_L(\bar{p},-p_\Delta,\lambda)-\pi_L(\underline{\mathbf{p}},\lambda)]-[\pi_H(\bar{p},p_\Delta,\lambda)-\pi_H(\underline{\mathbf{p}},\lambda)]\\
=&\left\{-\gamma\int_{-\infty}^{-\Delta-\frac{\lambda}{t}p_\Delta}F(\cdot)d\epsilon-h\left[\int_{-\infty}^{-\Delta-\frac{\lambda}{t}p_\Delta}F(\cdot)d\epsilon+\Delta\right]+\gamma\int_{-\infty}^{\Delta-\frac{\lambda}{t}p_\Delta}F(\cdot)d\epsilon+h\left[\int_{-\infty}^{\Delta-\frac{\lambda}{t}p_\Delta}F(\cdot)d\epsilon-\Delta\right]\right\}\\
&-\left\{-\gamma\int_{-\infty}^{-\Delta}F(\cdot)d\epsilon-h\left[\int_{-\infty}^{-\Delta}F(\cdot)d\epsilon+\Delta\right]+\gamma\int_{-\infty}^{\Delta}F(\cdot)d\epsilon+h\left[\int_{-\infty}^{\Delta}F(\cdot)d\epsilon-\Delta\right]\right\}\\
=&(\gamma+h)\left[\int_{-\Delta-\frac{\lambda}{t}p_\Delta}^{-\Delta}F(\cdot)d\epsilon-\int_{\Delta-\frac{\lambda}{t}p_\Delta}^{\Delta}F(\cdot)d\epsilon\right]<0
\end{aligned}$$

This means $[\pi_L^d(\underline{\mathbf{p}},\lambda)-\pi_L(\underline{\mathbf{p}},\lambda)]-[\pi_H^d(\underline{\mathbf{p}},\lambda)-\pi_H(\underline{\mathbf{p}},\lambda)]<0$ and Equation (23) is true. □

Lemma 14 establishes that from the collusive outcome $\mathbf{p}^c=\underline{\mathbf{p}}$, increasing p_H^c above $\underline{p}$ reduces buyer H's incentive to defect, which also shows existence of collusive outcomes with $\underline{p}<p_H^c\leqslant p_H^{c*},p_L^c=\underline{p}$.

Lemma 14. *If* $\Delta>0$, *given any collusive outcome* $\mathbf{p}^c$ *where* $p_H^c=\underline{p}+p_\Delta$ *and* $p_L^c=\underline{p}$, *we have* $\left.\frac{d\hat{\delta}_H(\mathbf{p}^c,\lambda)}{dp_\Delta}\right|_{p_\Delta=0}<0$ *for all* $\lambda>\max\{\hat{\lambda},\hat{\lambda}^c\}$.

Proof. It is straightforward to verify that for all $\lambda>\max\{\hat{\lambda},\hat{\lambda}^c\}$, $\pi_H^d(\mathbf{p}^c,\lambda),\bar{\pi}(\mathbf{p}^*(\lambda),\lambda)$ are independent of p_Δ, and $\left.\frac{d\pi_H(\mathbf{p}^c,\lambda)}{dp_\Delta}\right|_{p_\Delta=0}>0$, $\left.\frac{d\bar{\pi}(\mathbf{p}^c,\lambda)}{dp_\Delta}\right|_{p_\Delta=0}>0$. The result of the lemma then follows from the definition of $\hat{\delta}_H(\mathbf{p}^c,\lambda)$ given in Equation (22). □

Lemma 13 and 14 together imply that there always exist collusive equilibria $\mathbf{p}^c$ with $p_H^c(\lambda)>\underline{p}$ and $p_L^c(\lambda)=\underline{p}$ such that for any $\lambda>\max\{\hat{\lambda},\hat{\lambda}_c\}$, $\hat{\delta}(\mathbf{p}^c,\lambda)<\hat{\delta}(\underline{\mathbf{p}},\lambda)$. Combined with Lemma 12, we obtain the results in Theorem 3.

Appendix E: Proofs for Extensions and Additional Results

E.1. Proof of Proposition 7

All sellers strictly benefit for any $\lambda > \hat{\lambda}$ since $p_H^*(\lambda) > \underline{p}$ for all $\lambda > \hat{\lambda}$. For buyers, at $\lambda = \hat{\lambda}$, $\mathbf{p}^*(\lambda) = \underline{\mathbf{p}}$, and

$$\bar{\pi}(\mathbf{p}^*(\lambda), \hat{\lambda}) - \bar{\pi}(\underline{\mathbf{p}}, 0) = (\gamma + h) \int_{-\infty}^{-\Delta} [F(\epsilon; \sigma(0)) - F(\epsilon; \sigma(\hat{\lambda}))] d\epsilon > 0.$$

We used the fact that $\sigma(\hat{\lambda}) < \sigma(0)$. Since buyer's expected profit $\bar{\pi}(\mathbf{p}^*(\lambda), \lambda)$ is continuous in λ, we have the result in Proposition 7.

E.2. Proof of Theorem 4.A

Define $g_i(y) = \int_{-\infty}^{-|y-Q_i|} \frac{\partial F(\epsilon;\sigma(y))}{\partial y} d\epsilon = -\frac{\sigma'(y)}{\sigma(y)^2} \int_{-\infty}^{-|y-Q_i|} \epsilon f(\epsilon; \sigma(y)) d\epsilon$. Note that $\zeta = g_L(1) = g_H(1)$ and $g_i(y) \geqslant 0$. First, we establish the competitive market condition and that buyer H raises his price above $\hat{\lambda}$.

Lemma 15. *The competitive equilibrium prices satisfy $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ if and only if $\lambda \leqslant \hat{\lambda}$, where $\sigma_0 = \sigma(1)$, $\zeta = \int_{-\infty}^{-\Delta} \left[\frac{\partial F(\epsilon;\sigma(y))}{\partial y}\Big|_{y=1} \right] d\epsilon \geqslant 0$, and $\hat{\lambda} = \frac{t}{(\gamma+h)[F(\Delta;\sigma_0)-\zeta]-h+1-\underline{p}}$. Moreover, if $\Delta > 0$, there always exists $\lambda' > \hat{\lambda}$, such that for any $\lambda \in (\hat{\lambda}, \lambda')$, the competitive equilibrium prices satisfy $p_H^*(\lambda) > \underline{p}$ and $p_L^*(\lambda) = \underline{p}$, and $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) > \pi_H(\underline{p}, \underline{p}, 0)$.*

Proof. Differentiating buyer's individual expected profit with respect to price,

$$\frac{\partial \pi_H(p_H, p_L, \lambda)}{\partial p_H} = -2s_H - 2Q_H + 1 + \frac{\lambda}{t} \left[1 - p_L + \gamma - (\gamma + h)[F(s_H; \sigma(y_H)) + g_H(y_H)]\right].$$

$$\frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_L} = -2s_L - 2Q_L + 1 + \frac{\lambda}{t} \left[1 - p_H + \gamma - (\gamma + h)[F(s_L(\lambda); \sigma(y_L)) + g_L(y_L)]\right].$$

At $p_H = p_L = \underline{p}$, $s_L = -s_H = \Delta$. It is easy to verify that $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ if and only if $\lambda \leqslant \hat{\lambda}$ as both derivatives are $\leqslant 0$. When $\lambda > \hat{\lambda}$, $\frac{\partial \pi_H(p_H,p_L,\lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} > 0$ and thus $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ is no longer an equilibrium.

At $\hat{\lambda}$, if $\Delta > 0$, we have $\frac{\partial \pi_H(p_H,p_L,\lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} = 0$ and $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_L}\Big|_{p_H=p_L=\underline{p}} < 0$. By continuity, there exists $\lambda' > \hat{\lambda}$ such that for all $\lambda \in (\hat{\lambda}, \lambda')$, $p_H^*(\lambda) = \text{argmax}_{p_H} \pi_H(p_H, \underline{p}, \lambda) > \underline{p}$ and $p_L^*(\lambda) = \underline{p}$ because $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_L}\Big|_{p_H=p_H^*(\lambda),p_L=\underline{p}} \leqslant 0$. Since for all $\lambda \in (\hat{\lambda}, \lambda')$, $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) = \max_{p_H} \pi_H(p_H, \underline{p}, \lambda)$ and $\frac{\partial \pi_H(p_H,p_L,\lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} > 0$, we have $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) > \pi_H(\underline{p}, \underline{p}, \lambda) = \pi_H(\underline{p}, \underline{p}, 0)$ by optimality. □

Part 1, $\zeta \geqslant 0$. For general $\zeta \geqslant 0$, we show that $\Delta > \underline{\Delta}$, $h \geqslant 1 - \underline{p}$ leads to strong Pareto improvement. From Lemma 15, there exists $\lambda' > \hat{\lambda}$ such that all sellers and the high-demand buyer are strictly better off at $\lambda \in (\hat{\lambda}, \lambda')$. It remains to show that the low-demand buyer is *weakly* better off, such that both buyers are strictly better off when averaged over time.

Differentiating $\pi_L(p_L, p_H, \lambda)$ with respect to p_H, and using the shorthand $y_L = y(p_L, p_H, \lambda)$ and $s_L = \Delta - \frac{\lambda}{t}(p_H - p_L)$,

$$\frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_H} = -\frac{\lambda}{t} \left[1 - p_L + \gamma[1 - F(s_L; \sigma(y_L))] - hF(s_L; \sigma(y_L)) - (\gamma + h)g_L(y_L)\right].$$

Thus, $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_L=\underline{p}} \geqslant 0 \iff F(s_L; \sigma(y_L)) \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h} - g_L(y_L)$. let $\hat{p} = \underline{p} + \frac{t}{\lambda}(\Delta - \underline{\Delta}) > \underline{p}$. For any $p_H \in [\underline{p}, \hat{p}]$, $F(s_L; \sigma(y_L)) \geqslant F(\underline{\Delta}; \sigma(y_L)) \geqslant F(\underline{\Delta}; \sigma_0) \geqslant \frac{\gamma+1-\underline{p}}{\gamma+h} \geqslant \frac{1-\underline{p}+\gamma}{\gamma+h} - g_L(y_L)$. This implies $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_L=\underline{p}} \geqslant 0$ at all

$p_H \in [\underline{p}, \hat{p}]$. Hence, $\pi_L(\underline{p}, p_H, \lambda) = \pi_L(\underline{p}, \underline{p}, \lambda) + \int_{\underline{p}}^{p_H} \frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_H}\Big|_{p_L = \underline{p}} dp_H \geqslant \pi_L(\underline{p}, \underline{p}, \lambda)$ for all $p_H \in [\underline{p}, \hat{p}]$. This establishes that the low-demand buyer is at least weakly better off at any λ such that $p_H^*(\lambda) \in [\underline{p}, \hat{p}]$. According to Lemma 15, there exists $\lambda'' > \hat{\lambda}$ such that $p_H^*(\lambda) \in (\underline{p}, \hat{p}]$ for all $\lambda \in (\hat{\lambda}, \lambda'')$, where buyer H is strictly better off and buyer L is weakly better off. This leads to the first result of the theorem.

Part 2, $\zeta > 0$. We show that a tighter condition, $F(\Delta; \sigma_0) > \frac{1-\underline{p}+\gamma}{\gamma+h} - \zeta$ and $h + (\gamma + h)\zeta > 1 - \underline{p}$ leads to strong Pareto improvement. At $p_H = p_L = \underline{p}$, $F(\Delta; \sigma_0) > \frac{1-\underline{p}+\gamma}{\gamma+h} - \zeta$ implies that $\frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_H}\Big|_{p_H = p_L = \underline{p}} > 0$. By continuity, there must exist $\tilde{p} > \underline{p}$ such that $\pi_L(\underline{p}, p_H, \lambda) = \pi_L(\underline{p}, \underline{p}, \lambda) + \int_{\underline{p}}^{p_H} \frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_H}\Big|_{p_L = \underline{p}} dp_H \geqslant \pi_L(\underline{p}, \underline{p}, \lambda)$ for all $p_H \in [\underline{p}, \tilde{p}]$. Similarly, this establishes that the low-demand buyer is at least weakly better off at any λ such that $p_H^*(\lambda) \in [\underline{p}, \tilde{p}]$. According to Lemma 15, there exists $\lambda'' > \hat{\lambda}$ such that $p_H^*(\lambda) \in (\underline{p}, \tilde{p}]$ for all $\lambda \in (\hat{\lambda}, \lambda'')$, where buyer H is strictly better off and buyer L is weakly better off. This leads to the second result of the theorem.

E.3. Proof of Theorem 4.B

Before we present the general proof, we first show that when σ is a constant independent of y, the individual profit function of buyers remains concave. For the general case of $\sigma(y)$, we assume that concavity of the profit function is preserved (Assumption 3). Parallel to (P-FOC), for $h = 1 - p_i$, we have

$$\frac{\partial \pi_i(s_i, p_j, \lambda)}{\partial s_i} = -2s_i(\lambda) \cdot \frac{t}{\lambda} + 1 - p_j + \frac{t}{\lambda}(1 - 2Q_i) + \gamma[1 - F(s_i(\lambda); \sigma)] - (1 - p_i)F(s_i(\lambda); \sigma) + \frac{t}{\lambda}\int_{-\infty}^{s_i(\lambda)} F(\epsilon; \sigma) d\epsilon.$$

We verify that $\pi_i(s_i, p_j, \lambda)$ is concave by taking the second derivative,

$$\frac{\partial^2 \pi_i(s_i, p_j, \lambda)}{\partial s_i^2} = \begin{cases} -2 \cdot \frac{t}{\lambda} - (\gamma + 1 - p_i) f(s_i(\lambda); \sigma) + 2\frac{t}{\lambda} F(s_i(\lambda); \sigma) < 0 & \text{if } F(s_i(\lambda); \sigma) \in (0,1), \\ -2 \cdot \frac{t}{\lambda} + 2\frac{t}{\lambda} F(s_i(\lambda); \sigma) \leqslant 0 & \text{if } F(s_i(\lambda); \sigma) = 0 \, \text{or} \, 1. \end{cases}$$

Now, consider general $\sigma(y)$ that satisfies Assumption 3. Define $g_i(y) = \int_{-\infty}^{-|y-Q_i|} \frac{\partial F(\epsilon; \sigma(y))}{\partial y} d\epsilon = -\frac{\sigma'(y)}{\sigma(y)^2} \int_{-\infty}^{-|y-Q_i|} \epsilon f(\epsilon; \sigma(y)) d\epsilon$. Note that $\zeta = g_L(1) = g_H(1)$ and $g_i(y) \geqslant 0$. First, we establish that when $\Delta \geqslant \underline{\Delta}'$, buyers raise their price price if and only if $\hat{\lambda}$.

Lemma 16. *In any market where $\Delta \geqslant \underline{\Delta}'$, the competitive equilibrium prices satisfy $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ if and only if $\lambda \leqslant \hat{\lambda}$, where $\hat{\lambda} < 1$ and $\lambda > \hat{\lambda}$, where $\sigma_0 = \sigma(1)$, $\zeta = \int_{-\infty}^{-\Delta} \left[\frac{\partial F(\epsilon; \sigma(y))}{\partial y}\Big|_{y=1} \right] d\epsilon$, and $\hat{\lambda} = \frac{t(1 - \int_{-\infty}^{-\Delta} F(\epsilon; \sigma_0) d\epsilon)}{(\gamma + 1 - \underline{p})[F(\Delta; \sigma_0) - \zeta]}$. Moreover, there always exists $\lambda' > \hat{\lambda}$, such that for any $\lambda \in (\hat{\lambda}, \lambda')$, the competitive equilibrium prices satisfy $p_H^*(\lambda) > \underline{p}$ and $p_L^*(\lambda) = \underline{p}$, and $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) > \pi_H(\underline{p}, \underline{p}, 0)$.*

Proof. Differentiating buyer's individual expected profit with respect to price,

$$\frac{\partial \pi_H(p_H, p_L, \lambda)}{\partial p_H} = -2s_H - 2Q_H + 1 + \int_{-\infty}^{s_H} F(\epsilon; \sigma(y_H)) d\epsilon + \frac{\lambda}{t}\left[1 - p_L + \gamma - (\gamma + 1 - p_H)[F(s_H; \sigma(y_H)) + g_H(y_H)]\right].$$

$$\frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_L} = -2s_L - 2Q_L + 1 + \int_{-\infty}^{s_L} F(\epsilon; \sigma(y_L)) d\epsilon + \frac{\lambda}{t}\left[1 - p_H + \gamma - (\gamma + 1 - p_L)[F(s_L(\lambda); \sigma(y_L)) + g_L(y_L)]\right].$$

At $p_H = p_L = \underline{p}$, $s_L = -s_H = \Delta$. It is easy to verify that if $\Delta \geqslant \underline{\Delta}'$, $\frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_L}\Big|_{p_H = p_L = \underline{p}} < 0$ for all $\lambda \in (0, 1]$, and $\frac{\partial \pi_H(p_H, p_L, \lambda)}{\partial p_H}\Big|_{p_H = p_L = \underline{p}} \leqslant 0$ if and only if $\lambda \leqslant \hat{\lambda}$, where $\hat{\lambda} = \frac{t(1 - \int_{-\infty}^{\Delta} F(\epsilon; \sigma_0) d\epsilon)}{(\gamma + 1 - \underline{p})[F(\Delta; \sigma_0) - \zeta]}$.

Thus $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ if and only if $\lambda \leqslant \hat{\lambda}$ as both derivatives are $\leqslant 0$. When $\lambda > \hat{\lambda}$, $\frac{\partial \pi_H(p_H, p_L, \lambda)}{\partial p_H}\Big|_{p_H = p_L = \underline{p}} > 0$ and $p_H^*(\lambda) = p_L^*(\lambda) = \underline{p}$ is no longer an equilibrium.

At $\hat{\lambda}$, we have $\frac{\partial \pi_H(p_H,p_L,\lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} = 0$ and $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_L}\Big|_{p_H=p_L=\underline{p}} < 0$. By continuity, there exists $\lambda' > \hat{\lambda}$ such that for all $\lambda \in (\hat{\lambda}, \lambda')$, $p_H^*(\lambda) = \operatorname{argmax}_{p_H} \pi_H(p_H, \underline{p}, \lambda) > \underline{p}$ and $p_L^*(\lambda) = \underline{p}$ because $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_L}\Big|_{p_H=p_H^*(\lambda),p_L=\underline{p}} \leqslant 0$. Since for all $\lambda \in (\hat{\lambda}, \lambda')$, $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) = \max_{p_H} \pi_H(p_H, \underline{p}, \lambda)$ and $\frac{\partial \pi_H(p_H,p_L,\lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} > 0$, we have $\pi_H(p_H^*(\lambda), \underline{p}, \lambda) > \pi_H(\underline{p}, \underline{p}, \lambda) = \pi_H(\underline{p}, \underline{p}, 0)$ by optimality. □

To ensure that all sellers are in expectation better off, we also establish that there always exists $\lambda' > \hat{\lambda}$ such that for any $\lambda \in (\hat{\lambda}, \lambda')$, the expected amount of rejected supply is strictly less than the original market. As before, when $\lambda > \hat{\lambda}$, all sellers benefit from improved prices. We denote the expected amount of rejected supply each period using $w(p_H, p_L, \lambda)$. Then, $w(p_H, p_L, \lambda) = \int_{\infty}^{s_H(p_H,p_L,\lambda)} F(\epsilon; \sigma(y(p_H, p_L, \lambda)))d\epsilon + \int_{\infty}^{s_L(p_H,p_L,\lambda)} F(\epsilon; \sigma(y(p_L, p_H, \lambda)))d\epsilon$. At $p_H = p_L = \underline{p}$, w is strictly decreasing in p_H:

$$\frac{\partial w(p_H, p_L, \lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} = -\frac{\lambda}{t}\left[F(\Delta; \sigma_0) - F(-\Delta; \sigma_0) - \zeta + \zeta\right] < 0.$$

By continuity and Lemma 16, there always exists $\lambda' > \hat{\lambda}$ such that for any $\lambda \in (\hat{\lambda}, \lambda')$, the expected amount of rejected supply is strictly less than the original market. Thus, all the sellers are better off.

Next, we show that the buyers are better off.

Part 1, $\zeta \geqslant 0$. For general $\zeta \geqslant 0$, we show that $\Delta > \underline{\Delta}$ leads to strong Pareto improvement. From Lemma 15, there exists $\lambda' > \hat{\lambda}$ such that all sellers and the high-demand buyer are strictly better off at $\lambda \in (\hat{\lambda}, \lambda')$. It remains to show that the low-demand buyer is *weakly* better off, such that both buyers are strictly better off when averaged over time. Differentiating $\pi_L(p_L, p_H, \lambda)$ with respect to p_H, and using the shorthand $y_L = y(p_L, p_H, \lambda)$ and $s_L = \Delta - \frac{\lambda}{t}(p_H - p_L)$,

$$\frac{\partial \pi_L(p_L, p_H, \lambda)}{\partial p_H} = -\frac{\lambda}{t}\left[1 - p_L + \gamma[1 - F(s_L; \sigma(y_L))] - (1 - p_L)F(s_L; \sigma(y_L)) - (\gamma + 1 - p_L)g_L(y_L)\right].$$

Thus, $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_L=\underline{p}} \geqslant 0 \iff F(s_L; \sigma(y_L)) \geqslant 1 - g_L(y_L)$. let $\hat{p} = \underline{p} + \frac{t}{\lambda}(\Delta - \underline{\Delta}) > \underline{p}$. For any $p_H \in [\underline{p}, \hat{p}]$, $F(s_L; \sigma(y_L)) \geqslant F(\underline{\Delta}; \sigma(y_L)) \geqslant F(\underline{\Delta}; \sigma_0) = 1 \geqslant 1 - g_L(y_L)$. This implies $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_L=\underline{p}} \geqslant 0$ at all $p_H \in [\underline{p}, \hat{p}]$. Hence, $\pi_L(\underline{p}, p_H, \lambda) = \pi_L(\underline{p}, \underline{p}, \lambda) + \int_{\underline{p}}^{p_H} \frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_L=\underline{p}} dp_H \geqslant \pi_L(\underline{p}, \underline{p}, \lambda)$ for all $p_H \in [\underline{p}, \hat{p}]$. This establishes that the low-demand buyer is at least weakly better off at any λ such that $p_H^*(\lambda) \in [\underline{p}, \hat{p}]$. According to Lemma 16, there exists $\lambda'' > \hat{\lambda}$ such that $p_H^*(\lambda) \in (\underline{p}, \hat{p}]$ for all $\lambda \in (\hat{\lambda}, \lambda'')$, where buyer H is strictly better off and buyer L is weakly better off. This leads to the first result of the theorem.

Part 2, $\zeta > 0$. We show that a looser condition, $F(\Delta; \sigma_0) > 1 - \zeta$ leads to strong Pareto improvement. At $p_H = p_L = \underline{p}$, $F(\Delta; \sigma_0) > 1 - \zeta$ implies that $\frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_H=p_L=\underline{p}} > 0$. By continuity, there must exist $\tilde{p} > \underline{p}$ such that $\pi_L(\underline{p}, p_H, \lambda) = \pi_L(\underline{p}, \underline{p}, \lambda) + \int_{\underline{p}}^{p_H} \frac{\partial \pi_L(p_L,p_H,\lambda)}{\partial p_H}\Big|_{p_L=\underline{p}} dp_H \geqslant \pi_L(\underline{p}, \underline{p}, \lambda)$ for all $p_H \in [\underline{p}, \tilde{p}]$. Similarly, this establishes that the low-demand buyer is at least weakly better off at any λ such that $p_H^*(\lambda) \in [\underline{p}, \tilde{p}]$. According to Lemma 15, there exists $\lambda'' > \hat{\lambda}$ such that $p_H^*(\lambda) \in (\underline{p}, \tilde{p}]$ for all $\lambda \in (\hat{\lambda}, \lambda'')$, where buyer H is strictly better off and buyer L is weakly better off. This leads to the second result of the theorem.

E.4. Proof of Proposition 8

Under uniform reservation utility, the expected supply $y(p_i, p_j, \lambda)$ given in Equation (1) becomes

$$y'(p_i,p_j,\lambda)=\begin{cases}0 & \text{if } p_i<r,\\ 2(1-\lambda)\cdot\min\left\{\frac{1}{2},\frac{p_i-r}{t}\right\}+2\lambda\cdot\min\left\{\left(\frac{1}{2}+\frac{p_i-p_j}{2t}\right)^+,\frac{p_i-r}{t}\right\} & \text{if } r\leqslant p_i<r+t,\\ (1-\lambda)+2\lambda\cdot\min\left\{\left(\frac{1}{2}+\frac{p_i-p_j}{2t}\right)^+,1\right\} & \text{if } p_i\geqslant r+t.\end{cases}$$

In the initial market where $\lambda = 0$, the buyers' supplies are independent of each other, and the joint profit maximization problem is equivalent to individual profit maximization. Since sellers are loyal to the closer buyer, the optimal pricing must satisfy $p_L^{c*} \leqslant r + \frac{1}{2}t$, $p_H^{c*} \leqslant r + \frac{1}{2}t$, because increasing the price above $r + \frac{1}{2}t$ leads to no change in expected supply y' and lower expected profit. We use $p_L^{c*}(0) \leqslant r + \frac{1}{2}t$, $p_H^{c*}(0) \leqslant r + \frac{1}{2}t$ to denote profit-maximizing collusive price at $\lambda = 0$.

For any $\lambda > 0$, we first observe that the buyers will never be worse off compared to $\lambda = 0$ if they collude to maximize joint profit. This is because pricing at $p_L = p_L^{c*}(0)$, $p_H = p_H^{c*}(0)$ is always feasible and leads to the same expected profits as in the initial market: at $p_L^{c*}(0), p_H^{c*}(0)$, the buyers only capture their own loyal sellers regardless of information transparency λ, and so the market remains the same as the initial market. Thus jointly profit maximizing collusion can only lead to equal or higher expected profits for buyers.

It remains to show that the sellers will never be worse off compared to $\lambda = 0$. It is sufficient to show that $p_L^{c*}(\lambda) + p_H^{c*}(\lambda) \geqslant p_L^{c*}(0) + p_H^{c*}(0)$ for all $\lambda > 0$. We will show this result by establishing that among all pricing strategies that satisfy $p_L + p_H \leqslant p_L^{c*}(0) + p_H^{c*}(0)$, $p_L + p_H = p_L^{c*}(0) + p_H^{c*}(0)$ always leads to the highest expected profit, $\bar{\pi}$, under any $\lambda > 0$.

We consider three exhaustive cases that fall under $p_L + p_H \leqslant p_L^{c*}(0) + p_H^{c*}(0)$:

1. $r \leqslant p_L \leqslant p_L^{c*}(0)$ and $r \leqslant p_H \leqslant p_H^{c*}(0)$.
2. $r \leqslant p_L \leqslant p_L^{c*}(0)$, $p_H^{c*}(0) \leqslant p_H \leqslant r + t$, and $p_L + p_H \leqslant p_L^{c*}(0) + p_H^{c*}(0)$.
3. $r \leqslant p_H \leqslant p_H^{c*}(0)$, $p_L^{c*}(0) \leqslant p_L \leqslant r + t$, and $p_L + p_H \leqslant p_L^{c*}(0) + p_H^{c*}(0)$.

We show that the dominant pricing strategy must satisfy $p_L + p_H = p_L^{c*}(0) + p_H^{c*}(0)$ for all three cases for all $\lambda > 0$, and thus we must have $p_L^{c*}(\lambda) + p_H^{c*}(\lambda) \geqslant p_L^{c*}(0) + p_H^{c*}(0)$ for all $\lambda > 0$, which completes the proof for the proposition.

Case 1. $r \leqslant p_L \leqslant p_L^{c*}(0)$ and $r \leqslant p_H \leqslant p_H^{c*}(0)$. For this case, we now establish that, for any $\lambda > 0$, $p_L = p_L^{c*}(0), p_H = p_H^{c*}(0)$ gives the highest joint profit among all feasible pricing strategies. For any λ, whenever $r \leqslant p_L \leqslant r + \frac{1}{2}t$, $r \leqslant p_H \leqslant r + \frac{1}{2}t$, then the joint profit-maximization problem for both buyers can be decomposed to individual profit-maximization, and is independent of λ, because the buyers only capture their own loyal sellers. For buyer $i \in \{L, H\}$, we have

$$y'(p_i,p_j,\lambda)\equiv y'(p_i)=2\cdot\frac{p_i-r}{t},\quad s_i(p_i,p_j,\lambda)\equiv s_i(p_i)=y'(p_i)-Q_i,\quad \frac{dp_i}{ds_i}=\frac{t}{2},$$

and

$$L(s_i;\sigma)=\gamma\int_{-\infty}^{-s_i}F(\epsilon;\sigma)d\epsilon+h\left(\int_{-\infty}^{-s_i}F(\epsilon;\sigma)d\epsilon+s_i\right),$$
$$\pi_i(p_i,p_j,\lambda)=\pi_i(p_i)=y'(p_i)(1-p_i)-L(s_i(p_i);\sigma).$$

Differentiating,

$$\frac{\partial(\pi_i(p_i)+\pi_j(p_j))}{\partial p_i}=\frac{\partial \pi_i(p_i)}{\partial p_i}=\frac{2}{t}\left(1+r-2p_i+\gamma-(\gamma+h)\cdot F(s_i(p_i);\sigma)\right).$$

Note that the derivative above is monotonically decreasing in p_i, and thus this profit-maximization problem is strictly concave in p_i. Importantly, the value of the derivative is independent on transparency level, λ.

Thus, given initial profit-maximizing prices at $p_L^{c*}(0), p_H^{c*}(0) \leqslant r+\frac{1}{2}t$, we have that $p_i = p_i^{c*}(0)$ is still optimal for all $\lambda > 0$ in the regime of $r \leqslant p_L \leqslant p_L^{c*}(0)$ and $r \leqslant p_H \leqslant p_H^{c*}(0)$ since

$$\frac{\partial(\pi_i(p_i)+\pi_j(p_j))}{\partial p_i}=\frac{\partial \pi_i(p_i)}{\partial p_i}\geqslant 0 \quad \forall p_i \in (r, p_i^{c*}(0)), \forall \lambda.$$

We hence have that $p_L + p_H = p_L^{c*}(0) + p_H^{c*}(0)$ yields the highest expected profit among all pricing strategies that satisfy $r \leqslant p_L \leqslant p_L^{c*}(0)$ and $r \leqslant p_H \leqslant p_H^{c*}(0)$.

Case 2. $r \leqslant p_L \leqslant p_L^{c*}(0)$, $p_H^{c*}(0) \leqslant p_H \leqslant r+t$, and $p_L + p_H \leqslant p_L^{c*}(0) + p_H^{c*}(0)$. For this case, we now establish that for any $\lambda > 0$, among all feasible pricing strategies, the pricing strategy that gives the highest expected profit must satisfy $p_L + p_H = p_L^{c*}(0) + p_H^{c*}(0)$. First note that $p_H + p_L \leqslant p_L^{c*}(0) + p_H^{c*}(0)$ implies $p_H + p_L \leqslant 2r+t$, which in turn implies $\frac{p_H - r}{t} \leqslant \frac{1}{2} + \frac{p_H - p_L}{2t}$ and $\frac{p_L - r}{t} \leqslant \frac{1}{2} + \frac{p_L - p_H}{2t}$. Hence, the expressions for $y'(p_i, p_j, \lambda)$ in this case can be simplified as follows.

$$y'(p_H,p_L,\lambda)\equiv y'(p_H,\lambda)=\begin{cases}(1-\lambda)+2\lambda\cdot\frac{p_H-r}{t} & \text{if } p_H > r+\frac{1}{2}t\\ 2\cdot\frac{p_H-r}{t} & \text{if } p_H \leqslant r+\frac{1}{2}t\end{cases},$$

$$y'(p_L,p_H,\lambda)\equiv y'(p_L)=2\cdot\frac{p_L-r}{t}.$$

The joint profit-maximization problem for both buyers can again be decomposed to individual profit-maximization problems. We recognize that buyer L's profit maximization problem is identical to that under Case 1, and thus following the same proof we have, for any $\lambda > 0$,

$$\frac{\partial(\pi_L(p_L)+\pi_H(p_H))}{\partial p_L}=\frac{\partial \pi_L(p_L)}{\partial p_L}=\frac{2}{t}\left(1+r-2p_L+\gamma-(\gamma+h)\cdot F(s_L(p_L);\sigma)\right)\geqslant 0 \quad \forall p_L \in (r, p_L^{c*}(0)).$$

Hence for any λ, it is optimal for L to price as high as possible, i.e., $p_L = p_L^{c*}(0) + p_H^{c*}(0) - p_H$ is the optimal pricing strategy for buyer L. This suggests that the dominant pricing strategy in this case must satisfy $p_L + p_H = p_L^{c*}(0) + p_H^{c*}(0)$.

Case 3. $r \leqslant p_H \leqslant p_H^{c*}(0)$, $p_L^{c*}(0) \leqslant p_L \leqslant r+t$, and $p_L + p_H \leqslant p_L^{c*}(0) + p_H^{c*}(0)$. The proof is analogous to Case 2. We recognize that it is again optimal for H to price as high as possible since $\frac{\partial(\pi_L(p_L)+\pi_H(p_H))}{\partial p_H} \geqslant 0$. This suggests that the dominant pricing strategy must satisfy $p_L + p_H = p_L^{c*}(0) + p_H^{c*}(0)$.

E.5. Proof of Theorem 5

Let p_{LL}, p_{HH} denote prices when buyers are simultaneously in high and low demand, respectively. Let p_{HL}, p_{LH} denote the price of the high-demand buyer and the low-demand buyer when buyers are in opposite demand states, respectively. We start by characterizing the competitive equilibrium in Lemma 17.

Lemma 17. *The competitive equilibrium prices satisfy* $p^*_{LL}(\lambda) = p^*_{HH}(\lambda) = p^*_{HL}(\lambda) = p^*_{LH}(\lambda) = \underline{p}$ *for all* $\lambda \leqslant \hat{\lambda}$, *where* $\hat{\lambda} = \frac{t}{1-\underline{p}+\gamma-(\gamma+h)F(1-Q_H;\sigma)}$. *Moreover, if* $F(1-Q_L;\sigma) > F(1-Q_H;\sigma)$, *there always exists* $\lambda' > \hat{\lambda}$, *such that for any* $\lambda \in (\hat{\lambda}, \lambda')$, *the competitive equilibrium prices satisfy* $p^*_{HH}(\lambda) > \underline{p}$, $p^*_{HL}(\lambda) > \underline{p}$ *and* $p^*_{LL}(\lambda) = p^*_{LH}(\lambda) = \underline{p}$.

Proof. Differentiating buyer's individual expected profit with respect to price,

$$\frac{\partial \pi_i(p_i, p_j, \lambda)}{\partial p_i} = -2s_i(\lambda) - 2Q_i + 1 + \frac{\lambda}{t}\left[1 - p_j + \gamma - (\gamma+h)F(s_i(\lambda);\sigma)\right].$$

When buyers are both in the low demand state, $p_i = p_j = p^*_{LL}$. Given the FOC and the boundary condition that $p^*_{LL} \geqslant \underline{p}$, it is easy to verify that $p^*_{LL}(\lambda) = \max\left\{\underline{p}, 1 + \frac{\gamma-h}{2} - \frac{t}{\lambda} - (\gamma+h)\cdot\left[F(1-Q_L;\sigma) - \frac{1}{2}\right]\right\}$. Similarly, we have $p^*_{HH}(\lambda) = \max\left\{\underline{p}, 1 + \frac{\gamma-h}{2} - \frac{t}{\lambda} - (\gamma+h)\cdot\left[F(1-Q_H;\sigma) - \frac{1}{2}\right]\right\}$. p^*_{LL}, p^*_{HH} both satisfy the descriptions of the lemma.

When buyers are in opposite demand states, at $p_{HL} = p_{LH} = \underline{p}$, we have $\frac{\partial \pi_{HL}(p_{HL}, p_{LH}, \lambda)}{\partial p_{HL}} \leqslant 0$ for $\lambda \leqslant \hat{\lambda}$, and > 0 for $\lambda > \hat{\lambda}$. Meanwhile, $\frac{\partial \pi_{LH}(p_{LH}, p_{HL}, \lambda)}{\partial p_{LH}} < 0$ for $\lambda \leqslant \hat{\lambda}$. We have the result in Lemma 17 by continuity and concavity of buyer's profit function. □

Consequently, all sellers strictly improve when $\lambda > \hat{\lambda}$. It remains to show that buyers strictly benefit. Depending on the demand state of the market, buyer profits vary as follows:

1. Buyers are simultaneously in low demand states (w.p. $1 - 2\eta + \tau$). Based on Lemma 17, there exists $\lambda' > \lambda$ such that for any $\lambda \in (\hat{\lambda}, \lambda')$, buyer welfare remain unchanged compared to the original market, i.e., $\pi_{LL}(p^*_{LL}(\lambda), \lambda) = \pi_{LL}(\underline{p}, 0)$.
2. Buyers are in opposite demand states (w.p. $\eta - \tau$). Based on Lemma 17 and following similar proof steps to Theorem 2, there exists $\lambda' > \lambda$ such that for any $\lambda \in (\hat{\lambda}, \lambda')$, buyers strictly benefit compared to the original market, i.e., $\pi_{HL}(p^*_{HL}(\lambda), p^*_{LH}(\lambda), \lambda) + \pi_{LH}(p^*_{LH}(\lambda), p^*_{HL}(\lambda), \lambda) > \pi_{HL}(\underline{p}, \underline{p}, 0) + \pi_{LH}(\underline{p}, \underline{p}, 0)$.
3. Buyers are simultaneously in high demand states (w.p. τ). Buyers are always worse off for any $\lambda > \hat{\lambda}$, i.e., $\pi_{HH}(p^*_{HH}(\lambda), \lambda) < \pi_{HH}(\underline{p}, 0)$.

Thus, it is straightforward to verify that buyers are strictly better off for any $\lambda \in (\hat{\lambda}, \lambda')$ when averaged over all demand states, when τ is sufficiently small.