Heterogeneous Buyer Valuations and Binding Capacity Constraints Depress Contract Prices for Agricultural Products

Zhiran Qin[§] Richard J. Sexton[§]

July 2026

[§] Department of Agricultural and Resource Economics, University of California, Davis

Abstract

Contract procurement dominates today for most agricultural commodities in Western countries and is increasing in importance in the developing world. In the prototype setting, a few processors bid for the output of many farmers whose products they value heterogeneously based on location, brand characteristics, and other factors. Modern processing/packing facilities are capital intensive with fixed capacities. Despite the predominance of this market environment, it has largely been ignored in studies of farm price determination. We address this substantial gap in the literature by modeling contract bidding in this market environment, drawing on the theory of two-sided matching. Using the Hotelling model, we show that the joint presence of heterogeneous buyer valuations and binding capacity constraints can largely attenuate the bargaining power of farms and cause farm prices to be reduced to sellers' reservation or outside-option price even under non-collusive bidding. Any reduction in processor capacity within a market area weakly lowers every seller's contract price, such that limiting processing capacity can be an important strategic tool for buyers to depress farm prices.

Keywords: contract pricing, capacity constraints, heterogeneous valuation, two-sided matching, oligopsony power

1 Introduction

Key characteristics of modern agricultural markets are that (i) transactions between farmers and intermediaries (processors, shippers, retailers) for most commodities are governed by contracts (Goodhue, 2011; Wang, Wang, and Delgado, 2014; Crespi and MacDonald, 2022);[1] (ii) the intermediary stage of the supply chain is often highly concentrated within the relevant geographic markets (Rogers and Sexton, 1994; Deconinck, 2021; Crespi and MacDonald, 2022);[2] (iii) processors operate capital-intensive plants with fixed capacities (MacDonald, 2024; Moschini and Smith, 2026) and face the cost-based imperative to operate these plants at full capacity (Muth et al., 2005; U.S. Government Accountability Office, 2009); and (iv) farm products are differentiated based upon geographic location, objective quality, and/or presence of credence attributes pertaining to the production process—e.g., organic, non-GMO, local, sustainable, regenerative, animal friendly (Dries, Reardon, and Swinnen, 2004; Saitone and Sexton, 2017).

These differences create heterogeneous valuations among potential contract purchasers for any farm's products based on location of the farm relative to the buyer's processing/packing facility and the brands and product characteristics the buyer seeks to market to its downstream buyers. Most theory of price determination for agricultural products, however, is mired within models that assume spot-market exchange and homogeneous agricultural products. Such models either invoke perfect competition or assume simple models of oligopsony competition such as the flexible oligopoly-oligopsony market (FOOM) model, which reduces buyer power to a single parameter located on the unit interval (Sexton and Xia, 2018).[3] Studies of farm product procurement that depart from this convention focus typically upon the spatial dimension of farm-product procurement markets and invoke variants of the basic Hotelling model of price competition (Mérel, Sexton, and Suzuki, 2009; Fousekis, 2011; Graubner, Balmann, and Sexton, 2011; Graubner, Salhofer, and Tribl, 2021; Jung, Sesmero, and Siebert, 2022; Moschini and Smith, 2026).

These models, however, tell us little about how farm prices are set in modern agricultural markets characterized by the aforementioned four conditions—contract procurement

[1] Spot market exchange is limited in the United States to grains and oilseeds and as a residual market for other products including livestock, fruits, vegetables, and nuts (MacDonald and Korb, 2011). Contract farming is also becoming more pervasive in the developing world (Meemken and Bellemare, 2020).

[2] Because agricultural products are often bulky, perishable, and costly to transport, the relevant geographic markets for product procurement are local or regional in scope, meaning that concentration measures, such as the four-firm concentration ratio (CR4) or Hirshman-Herfindahl Index (HHI) computed at a national level, understate the concentration that exists within local or regional procurement markets (Sexton and Xia, 2018).

[3] This parameter, most often denoted as θ, takes on the value of zero for perfect competition, 1.0 for pure monopsony, with values $\theta \in (0, 1)$ denoting various degrees of oligopsony power.

wherein prices are determined through negotiations between buyers and sellers in settings where few buyers and relatively many sellers are present, and buyers are characterized by binding capacity constraints and heterogeneous valuations for any farm's production. Given that policymakers worldwide are intensely interested in farm prices and farmer welfare, understanding contract price determination in modern agricultural markets is imperative, and lending insight into this key issue is the focus of the present paper.

The prototype framework to study contract pricing for agricultural commodities is the principal-agent model (e.g., Hueth and Ligon, 1999; Goodhue, 2027). Such models cast the buyer as principal and farmers in the role of agents, wherein the power of the transaction is delegated to the principal, who, in essence, operates as a monopsony buyer. The buyer-principal selects contracts to maximize its payoff subject to meeting the farmer-agents' participation (individual rationality) and incentive-compatibility constraints. Competition enters the framework only indirectly through farmers' outside selling option that then determines the participation constraint. The outside option, however, is typically given as exogenous, with no discussion of the market processes or competitive environment that would determine such an option.

We develop a model of contract pricing that incorporates the key characteristics that define modern agricultural product procurement markets: (i) high buyer concentration and large numbers of farmer-sellers. Such markets are not perfectly competitive, but they are seldom monopsonies, as presumed implicitly by the principal-agent model. (ii) Buyers have heterogeneous valuations for any farm's products. (iii) Buyers have capacity constraints that may bind in market equilibrium (Jung, Sesmero, and Siebert, 2022; Moschini and Smith, 2026).

Even when the first two of these features are incorporated in models of farm product procurement (e.g., as with the spatial models cited previously), the third is almost always excluded.[4] Nonetheless, buyer capacity constraints are an essential feature of almost any procurement market. For example, Moschini and Smith (2026) find U.S. beef processing plants operating at an average of 96% of capacity in the baseline equilibrium, and that, across their counterfactual scenarios, fed cattle prices are most affected by the tightening or easing of capacity constraints.

In addition to processing/packing plants being constrained based upon their physical dimensions, other factors may also limit their ability to procure and handle farm product. One factor is the difficulty of expanding labor input due to either availability of workers to

[4]A recent exception is Moschini and Smith (2026), who incorporate "soft" capacity constraints into a calibrated spatial Bertrand model of U.S. beef packing. Capacity can be exceeded in their model at increasing marginal cost. In contrast, we model hard constraints that cannot be exceeded.

operate a second shift or the need to pay overtime wages to regular workers—a soft, cost-increasing constraint a la Moschini and Smith (2026) in the latter instance.[5] A second type of capacity constraint is that market intermediaries most often have sales contracts with downstream buyers that specify the specific volumes that a buyer agrees to acquire. Such contracts typically have a duration of a year or more, and are set in place well in advance of an intermediary's acquisition of the farm product input. The magnitude of farm product required to meet volumes contained in a buyer's sales contracts with its downstream partners forms its own constraint on product acquisitions quite apart from constraints imposed by physical capacities or labor availability.[6][7]

The failure of models of farm price determination to consider buyer capacity constraints is somewhat surprising, given that empirical studies have shown that binding processor capacity constraints are associated with significant markdowns of farm price. Recent examples for the U.S. include Jung, Sesmero, and Siebert (2022, 2026) for corn input to ethanol plants, Butcher and Schulz (2025) for hogs, and Moschini and Smith (2026) for fed cattle. Farmers also appear to understand the importance of buyer capacity to determining the prices they receive. A recent example involves announcements in 2026 of the closure of cattle processing plants by multiple of the "big four" cattle processors and the expressions of concern among ranchers and feedlot operators that followed.[8]

Our analytical framework draws on a literature that, to our knowledge, has not previously been applied to agricultural product procurement: the theory of two-sided matching markets with transferable utility, originating with work by Shapley and Shubik (1971) and extended by Crawford and Knoer (1981) and Kelso and Crawford (1982, hereafter KC) to the context of labor markets and firms' bidding to hire heterogeneous workers. Our analog is food processing firms with capacity to acquire product from multiple farms bidding contract

[5]Because processing facilities are often located in rural areas in order to be proximate to farm production, labor availability can represent an important factor constraining firms' production. Processing facilities have come to rely increasingly on immigrant labor, including undocumented migrants, e.g., Saitone, Schaefer, and Scheitrum (2021).

[6]Spot markets for selling a handler's finished products may simply not exist, or, if such selling opportunities do exist, they involve considerable price discounting, making the sum of contract sales volumes its own constraint on farm-product procurement.

[7]The industrial organization literature on price setting subject to capacity constraints has focused largely on sellers, not buyers, and emphasizes how capacity constraints attenuate price competition, with seminal contributions provided by Kreps and Scheinkman (1983) and Osborne and Pitchik (1986). A significant limiting feature of these models is that they assume sellers produce a homogeneous product and engage in price competition. Thus, absent binding capacity constraints, output price would be bid down to the level of seller marginal costs, i.e., Bertrand's paradox.

[8]Squire (2026) is just one among a great many popular-press articles reflecting this concern that appeared in the summer of 2026. Notably, in a competitive market environment, buyers shedding "excess capacity" and cutting costs could increase plant efficiency, reduce margins, increase farm prices, and thus not face farmer opposition.

prices for quantities of heterogeneous farm product.

We show in this environment that across a broad set of market conditions, equilibrium buyer contract offers succeed in depressing prices to the level of sellers' outside option or reservation price. We also isolate market conditions that enable farmers to earn higher returns. Our analysis of contract pricing in this setting proceeds in three parts. First, we illustrate the economic forces at work though some simple examples. They motivate the second component, a more general analysis based upon the theory of two-sided matching markets. This analysis establishes formally the conditions required for farm prices to collapse fully to sellers' outside option prices. Third, we apply the framework to the context of spatial competition between buyers based upon the Hotelling model (Hotelling, 1929), a key analytical framework used across three plus decades to study oligopsony competition in agricultural markets. We conclude with a discussion of implications of the analysis for farm prices and strategies and policies that can improve farm prices in the market environments we study.

2 Contract Bidding with Binding Capacities and Heterogeneous Valuations: Two Motivating Examples

We first illustrate the central logic of our framework with two examples of contract bidding between duopsony buyers. For concreteness, we often refer to the buyers as beef processors bidding for lots of cattle from feedlot-operator sellers, but the framework applies broadly to any farm-product procurement setting with high buyer concentration wherein buyers have heterogeneous valuations of available product and constraints on their capacity to procure and process it.

Suppose five feedlots $j = 1, \ldots, 5$ within a geographic market each supply one lot of cattle. Processors $i = 1, 2$ may submit bids on any available lot within the geographic market. Processor 1 has capacity for two lots and processor 2 has capacity for three, so available production exactly meets the sum of the processors' capacities.

A processor's valuation of a lot, V_{ij}, is the revenue attainable from selling output produced from the lot, less the marginal costs of transporting and processing it; valuations are heterogeneous across processors for the reasons given in the introduction. Each seller can exercise an outside option by selling its production outside the relevant market, e.g., for export, to a domestic processor located outside the geographic market, or to a residual spot market, if one exists. The value of the outside option is O_j, with $O_j < V_{1j}, V_{2j}$. Feedlots weakly prefer to sell into the local market and accept the higher local offer subject to it being

at least equal to O_j. Information is complete: each processor knows its own valuations and those of its rival.[9]

2.1 Example 1: High Valuations Match Capacities

Suppose first that processor 1 has the higher valuation for lots 1 and 2 and processor 2 for lots 3–5, so that each processor could exactly fill its capacity with the lots it values most. Table 1 summarizes the setup.

Table 1: The base example: two processors, five feedlots

	Processor 1		Processor 2	
Feedlot	Valuation	Guaranteed profit	Valuation	Guaranteed profit
1	V_{11}^*	$\theta_{11} = V_{11} - V_{21}$	V_{21}	—
2	V_{12}^*	$\theta_{12} = V_{12} - V_{22}$	V_{22}	—
3	V_{13}	—	V_{23}^*	$\theta_{23} = V_{23} - V_{13}$
4	V_{14}	—	V_{24}^*	$\theta_{24} = V_{24} - V_{14}$
5	V_{15}	—	V_{25}^*	$\theta_{25} = V_{25} - V_{15}$

Notes: An asterisk marks the high-valuation buyer for each lot. "Guaranteed profit" is the margin the high-valuation buyer earns even if forced to pay the rival's full valuation.

Each processor is guaranteed a minimum profit on the lots where it holds the higher valuation because the rival will never bid above its own valuation. By offering V_{21}, for example, processor 1 can obtain lot 1 because processor 2 will not bid above its valuation, earning at least $\theta_{11} = V_{11} - V_{21} > 0$. These margins appear in the "guaranteed profit" columns of Table 1. It is tempting to conclude that Table 1 also describes the market equilibrium: each winner pays a price equal to the second-highest bidder's valuation and earns the margin θ_{ij} as profit. Indeed this would be the Nash equilibrium under noncooperative processor bidding without capacity constraints, e.g., it is the equilibrium outcome of an English second-price auction. That conclusion is wrong, however. Equilibrium bids are uniformly *lower* in the presence of binding capacity constraints.

The conventional logic fails because it ignores what winning costs a capacity-constrained bidder in terms of foregone profits on its high-valuation lots. Following Table 1, let $\theta_1 = \min\{\theta_{11}, \theta_{12}\}$ and $\theta_2 = \min\{\theta_{23}, \theta_{24}, \theta_{25}\}$ denote the smallest guaranteed margin each processor earns per lot when it fills its capacity with its high-valuation lots; we call θ_i processor

[9]This assumption is consistent with the basic economics of food processing industries. Processors know the locations and capacities of rivals' processing facilities, the locations of available production, and the relevant transportation costs. They would also have good information on the types of final products each competitor seeks to produce and the quality characteristics of the available supplies of the farm product.

i's *profit floor*. It represents processor i's opportunity cost when it bids on other lots with a binding capacity constraint. Now consider processor 2 contemplating a bid on lot 1. Processor 2 has capacity for three lots, and it can fill all three with lots 3–5 at a per-lot margin of at least θ_2. Winning lot 1 would therefore force processor 2 to give up one of lots 3–5, sacrificing a margin of at least θ_2. It follows that processor 2 will never bid more than $V_{21}-\theta_2$ for lot 1; at any higher price, the profit that lot 1 offers is smaller than the margin on the lot it displaces. The same cap applies to lot 2, and, symmetrically, processor 1 never bids more than $V_{1j}-\theta_1$ on lots 3–5.

This simple observation illustrates an important point: bidders subject to capacity constraints will not bid up to their maximum valuations on lots for which they are not the high-valuation bidder. Those bids are constrained by the minimum profit that can be earned on lots wherein the firm is the high-valuation bidder and that can be used to reach its capacity constraint.

This, though, means that the high-valuation buyer need not and will not make a bid as high as the other buyer's valuation and, in turn, the minimum profit each processor can earn on lots where it has the high valuation is higher than indicated in Table 1 because the low-valuation processor will not bid up to its valuation. To secure lot 1, processor 1 need not outbid processor 2's full valuation V_{21}; outbidding processor 2's capped bid $V_{21}-\theta_2$ suffices. Processor 1's guaranteed margin on lot 1 therefore rises from θ_{11} to $\theta_{11}+\theta_2$, and likewise on lot 2. Table 2 records the updated profit guarantees. Each processor's profit floor has risen by its *rival's* profit floor: both floors now equal $\theta_1+\theta_2$, which we denote as Θ.

Table 2: The base example with revised bidding strategies

	Processor 1		Processor 2	
Feedlot	Valuation	Guaranteed profit	Valuation	Guaranteed profit
1	V_{11}^*	$\theta_{11}+\theta_2$	V_{21}	—
2	V_{12}^*	$\theta_{12}+\theta_2$	V_{22}	—
3	V_{13}	—	V_{23}^*	$\theta_{23}+\theta_1$
4	V_{14}	—	V_{24}^*	$\theta_{24}+\theta_1$
5	V_{15}	—	V_{25}^*	$\theta_{25}+\theta_1$

Notes: To win lot j, the high-valuation buyer now needs to outbid only the rival's capped bid $V_{kj}-\theta_k$, not the rival's full valuation V_{kj}; its guaranteed profit rises accordingly, from θ_{ij} in Table 1 to $\theta_{ij}+\theta_k$.

This same argument can be reapplied. The higher floors lower the bid ceilings once more to $V_{ij}-\Theta$, then $V_{ij}-\theta_i-\Theta$, and so on; every reduction in the rival's bid ceiling raises a processor's guaranteed margin by the same amount on lots for which it has the

high valuation. Each round of this iterative logic shifts another Θ per lot from prospective bids into guaranteed profits.[10] Assuming that all feedlots face the same outside option price O_j, pursuit of this logic to its ultimate conclusion generates the following equilibrium bid strategies: the high-valuation processor for each lot bids exactly the outside option O_j, and the rival bids below O_j or does not bid at all.

These strategies form a Nash equilibrium that is sustained by the buyers' capacity constraints. A bid below O_j cannot win, so O_j is the cheapest winning bid. Also, any bid greater than the feedlot's outside option would enable the non-high-valuation processor to obtain the lot, but, given the binding capacity constraint, this would require the processor to not acquire one of its high-valuation lots, which yields incremental profits $V_{ij} - O_j$, a greater profit than is obtainable from a non-high-valuation lot.[11] Equilibrium prices equal outside options on every lot, even though two well-informed, non-colluding buyers compete on price for every lot. We call this outcome, in which every contract price equals the seller's outside option, a *full collapse* of prices; when only some prices equal outside options and at least one exceeds it, we call the collapse *partial*. It is worth the reminder that, absent heterogeneous buyer valuations and binding capacity constraints, such price competition would raise price to the buyers' valuations, i.e., to the farm commodity's marginal value product, the Bertrand-Nash paradox solution.

The implication is that binding capacity constraints in the presence of heterogeneous buyer valuations can severely restrict contract bidding behavior. The result is not dependent on each processor being the high-valuation buyer on sufficient lots to satisfy its capacity constraint. We illustrate this point in the next subsection.

2.2 Example 2: High Valuations Do Not Match Capacities

Example 1 lacks generality in one key aspect; each processor could fill its capacity exactly with lots on which it holds the high valuation. Suppose instead that processor 1 has the high valuation only for lot 1, so that processor 2 is the high-valuation buyer for lots 2–5, four lots in total, thus exceeding its capacity by one unit. Table 3 summarizes the revised setup. Processor 1 must acquire a lot its rival values more in order to meet its capacity. An intuition that it must pay a premium to achieve this outcome is not correct. The capacity constraint does the same work as before, and prices again fall to the outside options.

[10] With both profit floors at Θ, processor k's bid cap on lot j falls to $V_{kj} - \Theta$, so the high-valuation processor's guaranteed margin rises to $\theta_{ij} + \Theta$ and its floor to $\theta_i + \Theta$; the next round's caps are $V_{ij} - \theta_i - \Theta$, and each further round lowers them again.

[11] This alignment between each processor's own ranking of lots and the division of high valuations is not innocuous; Section 4 elevates it to a general condition (there, *consistent rankings*) and shows that without it some prices exceed outside options.

Table 3: The extended example: processor 1 holds the high valuation only on lot 1

	Processor 1		Processor 2	
Feedlot	Valuation	Guaranteed profit	Valuation	Guaranteed profit
1	V_{11}^*	$\theta_{11} = V_{11} - V_{21}$	V_{21}	—
2	V_{12}	—	V_{22}^*	$\theta_{22} = V_{22} - V_{12}$
3	V_{13}	—	V_{23}^*	$\theta_{23} = V_{23} - V_{13}$
4	V_{14}	—	V_{24}^*	$\theta_{24} = V_{24} - V_{14}$
5	V_{15}	—	V_{25}^*	$\theta_{25} = V_{25} - V_{15}$

Notes: An asterisk marks the high-valuation buyer for each lot. Processor 2 now holds the high valuation on four lots, one more than it can process.

Processor 2 can only procure three of its four high-valuation lots, and it seeks to acquire the three carrying the largest margins. Let $\theta_2 = \min\{\theta_{22}, \theta_{23}, \theta_{24}, \theta_{25}\}$ and suppose, without loss of generality, that the smallest margin is on lot 2, so lot 2 is the lot processor 2 least wants among the four. Define $\tilde{\theta}_2 = \min\{\theta_{23}, \theta_{24}, \theta_{25}\}$ as the lowest profit margin on the three lots yielding the higher margins. By construction, $\tilde{\theta}_2 > \theta_{22}$.

Now apply the bid-ceiling principle of the base example to processor 2's bidding on lot 2. Winning lot 2 would displace one of lots 3–5 on which processor 2 earns a margin of at least $\tilde{\theta}_2$, so processor 2's maximum bid on lot 2 is $V_{22} - \tilde{\theta}_2$. Because $\tilde{\theta}_2 > \theta_{22}$, this maximum bid lies strictly below $V_{22} - \theta_{22} = V_{12}$, i.e., processor 1 can acquire lot 2 at a price strictly below its own valuation, earning a margin of at least $\tilde{\theta}_2 - \theta_{22} > 0$ on the lot despite it being the low-valuation buyer.

From here the "downward iteration" can be applied exactly as in the base example. The only amendment is that in this setting each processor's profit floor is the smallest margin on the lots it acquires in the equilibrium (lots 1 and 2 for processor 1, and lots 3–5 for processor 2), instead of on the lots where it happens to hold the high valuation. Processor 1's profit floor, $\min\{\theta_{11}, \tilde{\theta}_2 - \theta_{22}\}$, caps its bids on lots 3–5 below its valuations, and those caps raise processor 2's guaranteed margins on lots 3–5. The higher profits further reduce the maximum bid processor 2 is willing to make on lots 1 and 2, and each iteration of this logic again shifts profit from prospective bids into guaranteed profit floors for the processors.[12] With a uniform outside option, and the same ranking alignment as in the base example, the Nash equilibrium bids result in every contract price reduced to O_j.

The examples show that what disciplines bidding is not who values a lot the most,

[12]For example, processor 1's bid cap on lot 3 is $V_{13} - (\tilde{\theta}_2 - \theta_{22})$, which raises processor 2's guarantee on lot 3 to $\theta_{23} + (\tilde{\theta}_2 - \theta_{22})$ and in turn lowers its bid cap on lot 2 to $V_{22} - \theta_{23} - (\tilde{\theta}_2 - \theta_{22})$.

but the margin each processor earns on the marginal lot it acquires in equilibrium, given its capacity constraint. A processor may earn profit on a lot it values less than its rival, precisely because the rival's capacity commitments lower its bid caps. All feedlots faced the *same* outside option in the examples. We show in section 4 that this uniformity, together with capacities that just exhaust supply, comprises the mechanism that allows bids to be reduced to sellers' outside options.

The implications of the two examples are substantial. They involve non-collusive price competition between buyers. Yet, in the presence of heterogeneous buyer valuations and binding capacities, equilibrium farm prices are well below buyers' valuations of the products and may be no greater than sellers' outside option, i.e., sellers' participation constraints can be made to bind even though the bidding environment features multiple non-collusive buyers. The iterative logic, however, raises questions the examples cannot answer. What determines equilibrium prices under general configurations of valuations, capacities, and outside options? And how do prices respond when capacity changes? The next section seeks to answer these questions in a general setting.

3 General Results

The essential features of the contract bidding environment common to modern agricultural markets are consistent with the two-sided matching markets with transferable utility, as studied primarily in labor economics. Shapley and Shubik (1971) analyzed the one-to-one case, in which each buyer is matched to only one seller in the real estate market. Crawford and Knoer (1981) and Kelso and Crawford (1982) extended the theory to the labor-market environment and the many-to-one case, where firms with multi-worker demands bid salary offers to engage heterogeneous workers. That literature established when such markets possess equilibria, characterized the range of equilibrium salaries, and showed that which side of the market, buyers/firms or sellers/workers, proposes the terms determines where in that range prices settle. Hatfield and Milgrom (2005), whose notation we use in Appendix B to state the bidding process, later incorporated the KC labor market matching model as a specific case in their single framework of matching with contracts.

We draw on these results in this section to express the agricultural contract market in general form and establish what the general theory can tell us regarding this market environment. We keep formality to the minimum in this section, deferring technical details to appendices.

3.1 The Contract Market

The market involves n processor/buyers (e.g., beef processors) and m farmers/sellers (e.g., feedlots). Each feedlot supplies one lot of cattle and sells to at most one processor. As in Section 2, processor i's valuation of lot j viewed in isolation, without consideration of capacity constraints, is $V_{ij} > 0$, and each feedlot has an outside option $O_j > 0$. A feedlot's payoff is the price it receives, or O_j if it elects the outside option; feedlots accept the highest offer weakly above O_j. As in Section 2, all valuations and outside options are common knowledge.

Processor i faces a capacity constraint that enables it to acquire and process at most q_i lots. A *price vector* $\boldsymbol{p} = (p_1, \ldots, p_m)$ assigns one price to each lot. Processor i's profit from acquiring the set of lots $C_i \subseteq \{1, \ldots, m\}$ at the price vector $\boldsymbol{p}$ is

$$\pi^i(C_i;\, \boldsymbol{p}) \;=\; \sum_{j \in C} \big(V_{ij} - p_j\big), \qquad |C_i| \leq q_i. \tag{1}$$

We rank the lots by the margin they offer processor i: $V_{ij} - p_j$. Processor i demands the q_i highest-ranked lots among those with positive margins, and fewer than q_i only if fewer than q_i lots have positive margins. This set of lots is processor i's *demand* at price vector $\boldsymbol{p}$.

One instance where demand plays a central role is in Section 4. We write S_i for processor i's *reservation demand* when every lot is priced at its outside option $\boldsymbol{O} = (O_1, \ldots, O_m)$. It represents processor i's most favorable bidding choices, the (up to) q_i lots on which the margin $V_{ij} - O_j$ is highest. Ties at reservation prices are a matter of primitives, and we rule them out, so S_i is a single set of lots.

3.2 Nash Equilibrium: Characterization and Multiplicity

An *outcome* of the market refers to the assignment of pairwise disjoint sets of lots $C_1, \ldots, C_n$ to the n processors, with $|C_i| \leq q_i$, together with a price for each lot sold to a processor. Every feedlot excluded from $C_1, \ldots, C_n$ is sold at its outside option. Outcomes arise from a two-stage bidding game. In the first stage, processors simultaneously submit personalized bids. Processor i's strategy is a bid vector $\boldsymbol{b}_i = (b_{i1}, \ldots, b_{im}) \in \mathbb{R}^m_+$, with a price b_{ij} on each lot it pursues and $b_{ij} = 0$ on each lot on which it makes no offer. In the second stage, each feedlot accepts the highest bid if that bid is weakly above its outside option, and otherwise exercises its outside option. Feedlot j's price is therefore $p_j = \max\{O_j, \max_i b_{ij}\}$. Since $O_j > 0$, a bid below O_j is never accepted, so $b_{ij} = 0$ is equivalent to making no offer. Each processor pays its own bid on every lot it wins, so processor i winning the set of lots C_i earns $\pi^i(C_i;\, \boldsymbol{b}_i) = \sum_{j \in C_i}(V_{ij} - b_{ij})$.

The price vector $\boldsymbol{p}$ therefore collects the best offers across all processors and coincides with no single processor's bid vector. A *Nash equilibrium* of the bidding game is a profile of bid vectors $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ at which no processor can raise its profit by changing its own bids, given its rivals' bids and the feedlots' acceptance rule. Each Nash equilibrium yields an outcome that we call a *non-collusive equilibrium*: the allocation of lots together with a price vector $\boldsymbol{p}$ that records the executed contract price of every lot sold to a processor and the outside option O_j of every lot sold at the outside option.

Non-collusive equilibria admit a simple characterization. Appendix A.2 shows that an outcome is a non-collusive equilibrium if and only if three conditions hold: (i) Every feedlot that sells to a processor receives at least its outside option. (ii) Every processor's set of lots is a most profitable set at $\boldsymbol{p}$: no processor can raise its profit at $\boldsymbol{p}$ by exchanging a lot it holds for one it does not, by dropping a lot, or by adding a lot within its capacity. (iii) Every lot sold at the outside option is priced at its outside option in $\boldsymbol{p}$. The last two conditions say that the market clears: no processor wants a lot it does not hold at $\boldsymbol{p}$, and no lot that receives a bid at or above its outside option from a processor sells outside.

More than one Nash equilibrium (and the corresponding outcome) can exist for the market, and the base example of Section 2.1 can illustrate this result. Consider first the outcome the English second-price auction intuition predicts: each winner pays a price equal to its rival's valuation. Both processors bid $V_{21}, V_{22}, V_{13}, V_{14}, V_{15}$ for lots 1–5 in the end, and processor 1 wins on lots 1–2 and processor 2 wins on lots 3–5. This is a Nash equilibrium. At these prices, all feedlots receive the same bidding prices from the two processors, and there is no incentive to deviate from the contract. Processor 1 earns positive profits on lots 1–2 and exactly zero on lots 3–5, so it demands precisely $\{1, 2\}$ and won't deviate. Processor 2 is in the symmetric position.

On the other hand, the full-collapse of prices to feedlots' outside options described in Section 2.1 is a Nash equilibrium as well: each processor holds its q_i best lots, and neither can poach a rival's lot without giving up a lot that pays at least its profit floor. The two equilibria differ only in the division of surplus: the same lots go to the same processors, but every feedlot is paid the competing processor's valuation in one and its outside option in the other. This multiplicity raises the question of which equilibrium the game settles at, and in particular whether the answer favors processors or feedlots. The next subsection shows that settings wherein processors, not sellers, make bid offers results in the Nash equilibrium that favors processors.

3.3 Which Equilibrium? Bidding Favors the Proposing Side

Existence of multiple Nash equilibria gives a range of possible contract prices. Which equilibrium prevails depends on which side of the market proposes the contract terms. The nearly universal rule in agricultural contract procurement is that the processor-buyers propose the terms in the first stage and farmers-sellers accept the highest offer in the second stage. As long as bids are acceptable to sellers, processors can offer the ones most favorable to themselves.

We call the equilibrium with the lowest purchase prices *buyer-optimal*, where every processor earns at least as much as at any other Nash equilibrium. This subsection argues that in this bidding environment the market settles at the buyer-optimal equilibrium. We first show this in the base example by discarding bids no processor would like to offer, and then state the general result as Proposition 1.

The two Nash equilibria of the base example discussed in Section 3.2 are not equally credible. The high prices of the English-auction equilibrium are sustained by bids the losing processor would regret winning. Recall that $\theta_1 = \min\{V_{11} - V_{21},\, V_{12} - V_{22}\}$ is the smallest per-lot profit processor 1 secures on lots 1 and 2. A bid $b_{1j} > V_{1j} - \theta_1$ on a lot $j \in \{3, 4, 5\}$ is a *regret bid*: if accepted, it displaces a lot worth at least θ_1 in exchange for a profit below θ_1; if rejected, it changes nothing for processor 1 and only affects the price processor 2 pays. Thus, the profit on a regret bid is zero or negative. Bidding the full valuations V_{13}, V_{14}, V_{15}, as the English-auction equilibrium requires, is a regret bid on every lot beyond processor 1's capacity, and the same holds for processor 2's bids V_{21}, V_{22} on lots 1 and 2.

To formally rule out equilibria sustained by such regret bids, we apply the concept of *iterated admissibility*, the successive elimination of weakly dominated strategies (Moulin, 1979; Marx and Swinkels, 1997; Brandenburger, Friedenberg, and Keisler, 2008), and retain only the Nash equilibria with bids that survive every round of elimination. Note that these "rounds" are steps of reasoning or logic rather than a sequence of play: actual bids are only submitted once in stage 1.

In the first round, bids above valuation are weakly dominated for both processors, which is what makes θ_1 and θ_2 guaranteed floors. In the second round, given that processor 2 never bids above its valuations and that processor 1 holds lots 1 and 2 in every equilibrium outcome, regret bids are weakly dominated for processor 1: bidding at the cap $V_{1j} - \theta_1$ does at least as well against every remaining rival strategy and strictly better whenever a higher bid would be accepted. Under complete information, processor 2 anticipates this and needs to bid no more than $V_{1j} - \theta_1$ on lots 3–5, which raises its own floor on those lots to $\theta_2 + \theta_1$ and lowers the cap on its bids for lots 1 and 2 to $V_{2j} - \theta_2 - \theta_1$. Each further round lowers the surviving caps by at least $\min\{\theta_1, \theta_2\}$, so every bid on a rival's lot reaches the outside option

in finitely many rounds of this bidding logic whenever $\theta_1, \theta_2 > 0$. The only equilibrium that survives is the one in which both processors bid outside options, which is the buyer-optimal equilibrium of the base example.

That every lot is priced at its outside option is a feature of this example, in which neither processor's bid caps on its rival's lots ends up above the outside option. It does not have to be the case generally. In a general market, the processor-proposing aspect of the bidding process concerns the selection among multiple equilibria and guarantees that the equilibrium is buyer-optimal.

Proposition 1 (Existence and buyer-optimal selection)**.** *In the contract market:*

(i) A pure-strategy Nash equilibrium of the bidding game exists.

(ii) There is a unique buyer-optimal Nash equilibrium outcome, and at it every seller receives the lowest price it obtains at any Nash equilibrium.

(iii) When processors/buyers propose the terms, contract bidding settles at the buyer-optimal outcome.

The proof is in Appendix C.

Part (i) claims at least one Nash equilibrium exists. Part (ii) locates the equilibrium that matters under buyer contract offers. As illustrated in the aforementioned example, equilibrium prices form a range. Toward the high end of the range lies the English-auction outcome of Section 3.2, in which each feedlot is paid the second-highest processor's valuation. At the bottom of the price range, each feedlot receives the lowest price at which no rival processor demands its lot. Part (ii) says that the bottom is well defined. Moreover, since the interests of the two sides are perfectly opposed, as in two-sided matching markets generally (Shapley and Shubik, 1971; Roth and Sotomayor, 1990), the bottom equilibrium that maximizes every processor's profit is also the one that minimizes the price every feedlot receives.

Part (iii) is the central result. It states in general form what the examples showed: whatever division of surplus the market can support, a structure wherein processors offer bids delivers the outcome most favorable to processors. Note that Proposition 1 says nothing about outside options. The buyer-optimal price of a feedlot equals its outside option only when no rival processor is willing to pay more than the outside option for it, and Section 4 characterizes exactly when that is the case.

Existence in part (i) is not automatic. Moschini and Smith (2026) note that when producers simply sell to the highest offer, price competition among capacity-constrained processors is difficult to characterize and may require mixed strategies. Therefore, they add logit preference shocks to producers' choice of processor in their calibrated model of U.S.

beef packing industry. Our part (i) shows that in the contract market this highest-offer case has a pure-strategy equilibrium when processors have a hard capacity constraint that cannot be exceeded.

4 How Far Do Prices Fall? Consistent Rankings and Sellers' Outside Options

Every feedlot's price is bounded below by its outside option, so the buyer-optimal prices of Proposition 1 lie somewhere between O_j and the top of the equilibrium range. The examples of Section 2 display cases where every seller is paid exactly its outside option. However, full collapse to the outside-option prices is not always the result. In equilibrium, let λ_i be the smallest margin $V_{ij} - p_j$ that processor i earns on the lots it buys, or zero if its capacity does not bind. Acquiring another lot l would cost processor i at least λ_i, so it bids at most $V_{il} - \lambda_i$ for that lot. When this maximum bid exceeds the outside option O_l for some lot it does not buy, contract prices do not collapse fully to outside options.[13]

This section gives the exact condition on primitives under which that happens. The condition is of independent interest because its failure modes are economically meaningful. They identify settings when sellers can earn a premium above their outside options even at the buyer-optimal equilibrium, and why.

Specifically, the reservation demand in Section 3.1 matters here. Setting every seller's price to its outside option, processor i's reservation demand S_i represents the lots it intends to buy: the (up to) q_i lots on which its margin $V_{ij} - O_j$ is highest. Two cases arise thereafter. In one case no seller appears on more than one processor's reservation demand S_i. In this setting, each processor can obtain exactly the lots it most wants, and no processor i is outbid by another for the lots in S_i. At the buyer-optimal Nash equilibrium, every contract price then equals the seller's outside option.

Suppose instead that some seller's production appears in the reservation demands of two processors. Those processors hence have to bid against each other for it, and its price is pushed above the outside option. The condition for full collapse of prices to outside options is therefore that processors' reservation demands do not overlap:

[13] A two-processor, two-lot example can illustrate this point. Let $q_1 = q_2 = 1$ (one-lot capacities), $O_1 = O_2 = 1$, $V_{11} = 10$, $V_{12} = 3$, $V_{21} = 9$, and $V_{22} = 4$. One can verify that $(p_1, p_2) = (6, 1)$ are the buyer-optimal equilibrium prices, with processor 1 buying lot 1 and processor 2 buying lot 2. As in Section 2.1, each processor is the high-valuation buyer on one lot and capacities exactly exhaust supply. But the ranking alignment assumed there does not apply here: processor 2 values lot 1 above lot 2. Processor 1's smallest margin is $\lambda_1 = 10 - 6 = 4$, so its maximum bid for lot 2 is $V_{12} - \lambda_1 = -1 < O_2$, and lot 2 sells at $p_2 = O_2 = 1$. Processor 2's smallest margin is then $\lambda_2 = 4 - 1 = 3$, and its maximum bid for lot 1 is $V_{21} - \lambda_2 = 6$, which processor 1 must match. Therefore, lot 1 sells at $p_1 = 6 > O_1$, above the outside option.

Definition 1 (Consistent rankings, CR)**.** Processors' rankings are *consistent* if their reservation demands $S_1, \ldots, S_n$ are pairwise disjoint.

CR is a property of the market's primitives that concerns how each processor ranks lots. Every feedlot allocated to processor i in the proposed equilibrium must be among processor i's own q_i best procurement opportunities when every lot is priced at its outside option.

Two consequences are worth recording here. First, suppose every processor values every lot above its outside option. Then a reservation demand uses the processor's full capacity, and disjointness requires $\sum_i q_i \leq m$; CR requires the absence of excess processor capacity. Second, if in addition the outside options are uniform, CR is a condition on valuations alone: the processors' sets of q_i highest-valued lots must not overlap. In Example 1 (Section 2.1), these sets are exactly the lots on which each processor is the high-valuation buyer. Example 2 (Section 2.2) shows that CR can hold even when they are not: processor 1's two highest-valued lots are lots 1 and 2, although it is the high-valuation buyer only on lot 1.

Proposition 2 (Full collapse of prices to sellers' outside options if and only if consistent rankings applies)**.** *Every contract price at the buyer-optimal Nash equilibrium equals the seller's outside option if and only if the CR condition holds.*

The proof is in Appendix C. When CR holds, the allocation in which each processor i buys the lots in S_i, together with the prices $p = \boldsymbol{O}$, is the buyer-optimal equilibrium. All processors' reservation demands do not overlap, so each processor i simply offers O_j on the lots in S_i, and no seller receives a competing offer. When CR fails, at least one seller is paid strictly more than its outside option at every Nash equilibrium, including the buyer-optimal one.

Proposition 2 states in general form the mechanism illustrated in Section 2. A capacity-constrained processor's smallest margin λ_i is taken over the lots it acquires at the buyer-optimal equilibrium. Winning any other lot would force it to give up one of these lots, so it bids for that lot only up to $V_{ij} - \lambda_i$. CR is exactly the condition under which no such bid exceeds the lot's outside option, so every lot sells at its outside option. When CR fails, there is only a partial collapse of prices to outside options, though the outcome is still the buyer-optimal equilibrium by Proposition 1.

Failure of CR to apply can happen in two situations. The first is excess buyer capacity, as illustrated by Figure 2 in section 5. Reservation demands then overlap. Some processor has unused capacity, so $\lambda_i = 0$ and it bids up to its full valuation for additional lots. Prices then exceed outside options on the lots that appear in two processors' reservation

demands. The second is misaligned rankings, where the same lot appears among the q_i highest margins $V_{ij} - O_j$ of two different processors even when capacity is fully committed. The two-processor example in footnote 13 shows that valuations alone can produce the misalignment. Misaligned rankings can also result from heterogeneous outside options, which re-rank lots relative to valuations. Section 5 develops both situations in the spatial model.

We close this section with the comparative statics of capacity. Binding capacity is what mutes bidding, so it is natural to ask how prices respond when capacity changes.

Proposition 3 (Capacity comparative statics)**.** *At the buyer-optimal Nash equilibrium, expansion of any processor's capacity weakly raises every seller's contract price, and any reduction weakly lowers it.*

The proof is in Appendix C.

The direction is intuitive. Extra capacity lowers the expanding processor's smallest margin λ_i, which raises its bids on lots it does not hold, and higher bids can only help sellers. What the general framework adds is the strength of the statement. The result holds feedlot by feedlot for the entire market and for any capacity change: a plant closure, restricted operating hours, or a credible cut in processing speed.

Proposition 3 places the strategic-capacity discussion of Section 6 on a formal footing. Capacity is more than a simple physical constraint on processors. It is an instrument that can be reduced strategically to shift price against sellers.

5 The Hotelling Model with Processor Capacity Constraints

As noted, the Hotelling model (Hotelling, 1929) has provided a workhorse framework to study spatial oligopsony competition to acquire agricultural products (Faminow and Benson, 1990; Mérel, Sexton, and Suzuki, 2009; Graubner, Balmann, and Sexton, 2011; Fousekis, 2011; Sesmero, Balagtas, and Pratt, 2015; Graubner, Salhofer, and Tribl, 2021; Evans and Sesmero, 2024). Most applications have involved duopsony competition between processors located at the endpoints of Hotelling's line, with length normalized to 1.0. Heterogeneity in processor valuations emerges naturally amidst the spatial distribution of farms and processors with costly transportation.

To date, however, the applications of Hotelling to agricultural product procurement involve spot-market models that seek a Bertrand-Nash equilibrium in uniform price offers for each processor, with farmers choosing the buyer offering the higher net price. Specific results regarding prices received are not easily summarized due to differing assumptions

in the various papers regarding who bears the transport costs—buyers, sellers, or in some combination,[14] whether one of the buyers is a producer cooperative as in Fousekis (2011), and whether the supply of farm products is elastic or inelastic, with an implicit outside option price of zero.

Despite the breadth in modeling frameworks that this literature exhibits, none of the papers have considered the role that processors' capacity constraints play in determination of farm prices. We show here that introducing this realistic feature into the Hotelling model and allowing personalized contract bids has stark implications for farm prices, given that heterogeneous buyer valuations are already present as in integral feature of the Hotelling model.

Producers or feedlots are located continuously along Hotelling's line. Processors, located at the market's endpoints, procure cattle at the feedlot gate and transport them to the processing facility.[15] Heterogeneous buyer valuations emerge due to costly transportation of cattle. Processor 1 has the higher valuation for lots located closer to its plant, and likewise for processor 2. Processor 1, located at the left endpoint, has valuation of a lot of cattle at location x_j of $V_1(x_j) = \bar{V}_1 - tx_j$. Processor 2's valuation is $V_2(x_j) = \bar{V}_2 - t(1 - x_j)$, where t is the transport cost per unit distance and $\bar{V}_i$ is the value of a lot at the plant gate. For simplicity we assume $\bar{V}_1 = \bar{V}_2 = \bar{V}$, so valuations are equal at the market midpoint $x = 0.5$. In lieu of selling to either processor 1 or 2, each feedlot also has an outside option $O(x_j) = O$ for all j. The model is a continuum analog of the contract market of Section 3, and Propositions 1–3 apply to this setting.

5.1 Binding Capacity Constraints

We first consider the case where total feedlot production in the market area exceeds the processors' combined total capacity. Without loss of generality, we assume processor 1 has the smaller capacity, $\bar{Q}_1$, with processor 2's capacity $1 - \bar{Q}_2$ such that $\bar{Q}_1 < \bar{Q}_2$. The

[14] Zhang and Sexton (2001) compare market outcomes under an FOB price (farmers bear transport cost) or uniform-delivered price (seller bears transport cost), while Graubner, Balmann, and Sexton (2011) study settings when transport costs are shared between buyers and sellers. Although prices differ across these market environments, in all cases farmers receive the same price.

[15] Thus, we model directly the spatial environment where processors arrange transportation and bear the cost. Processors in general will be more efficient at transportation than farmers, so processor-arranged transportation is the norm in the U.S. and most Western economies. Processors may still charge farmers a delivery fee in these settings. In the specific case of processor-feedlot contracts in the U.S., the processor arranges and pays for transportation under some but not all contract forms. Our analysis would also apply to the FOB pricing case wherein farmers bear the transport cost, although the model would need to be recast. Specifically, buyers would have homogeneous valuations of product delivered to the plant gate, but buyers' contract bids would need to account for the heterogeneity in farmers' costs of transporting product to processor 1 vs. processor 2, e.g., processor 1 would need to outbid processor 2 to obtain production of a feedlot located closer to processor 2.

consistent rankings (CR) condition from section 4 holds in this market environment.

$V_1(x_j)$ is decreasing and $V_2(x_j)$ is increasing in a seller's location, x_j along Hotelling's line. Each processor's reservation demand S_i is an interval starting at its own endpoint of the line, $[0, \bar{Q}_1]$ and $[\bar{Q}_2, 1]$. These intervals are disjoint if and only if total capacity does not exceed total supply. Because transport costs rise with distance, each processor ranks the lots nearest its own plant highest, so CR holds and Proposition 2 applies, resulting in the model delivering the full collapse of equilibrium price to the sellers' uniform outside option without further conditions.

Figure 1 illustrates the market forces that diminish bids below processors' valuations. Processor 1's profit floor is $\theta_1 = V_1(\bar{Q}_1) - V_2(\bar{Q}_1)$, the margin on its marginal lot. Because of its capacity constraint, processor 1 will not bid its valuation on lots to the right of $\bar{Q}_1$ even where it is the high-valuation buyer (the interval $[\bar{Q}_1, Q^*]$ in figure 1); its maximum bid there is $V_1(x_j) - \theta_1$, the red line. This caps the highest bid processor 2 faces from processor 1, guaranteeing processor 2 a profit floor $\hat{\theta}_2 = V_2(\bar{Q}_2) - (V_1(\bar{Q}_2) - \theta_1)$ on its marginal lot, so processor 2 bids no higher than $V_2(x_j) - \hat{\theta}_2$ on the interval $[0, \bar{Q}_2]$, the blue line. Each round of this logic lowers both processors' maximum bids (the dashed lines in Figure 1), and the iteration converges to Nash equilibrium bids equal to the uniform outside option $O(x_j)$ for every feedlot.[16] The shaded areas represent the processors' profits.

This full collapse of farm prices to the outside option deserves emphasis. The two processors do not collude, and their procurement areas overlap, but when total production exceeds total processing capacity, their competition does nothing to raise prices. Contract bidding, binding capacity constraints, and heterogeneous valuations together break the English second-price auction prediction that the high-valuation buyer wins the lot but pays the rival's valuation. Each processor therefore acts as a local monopsonist and secures its supply by offering just the outside option.

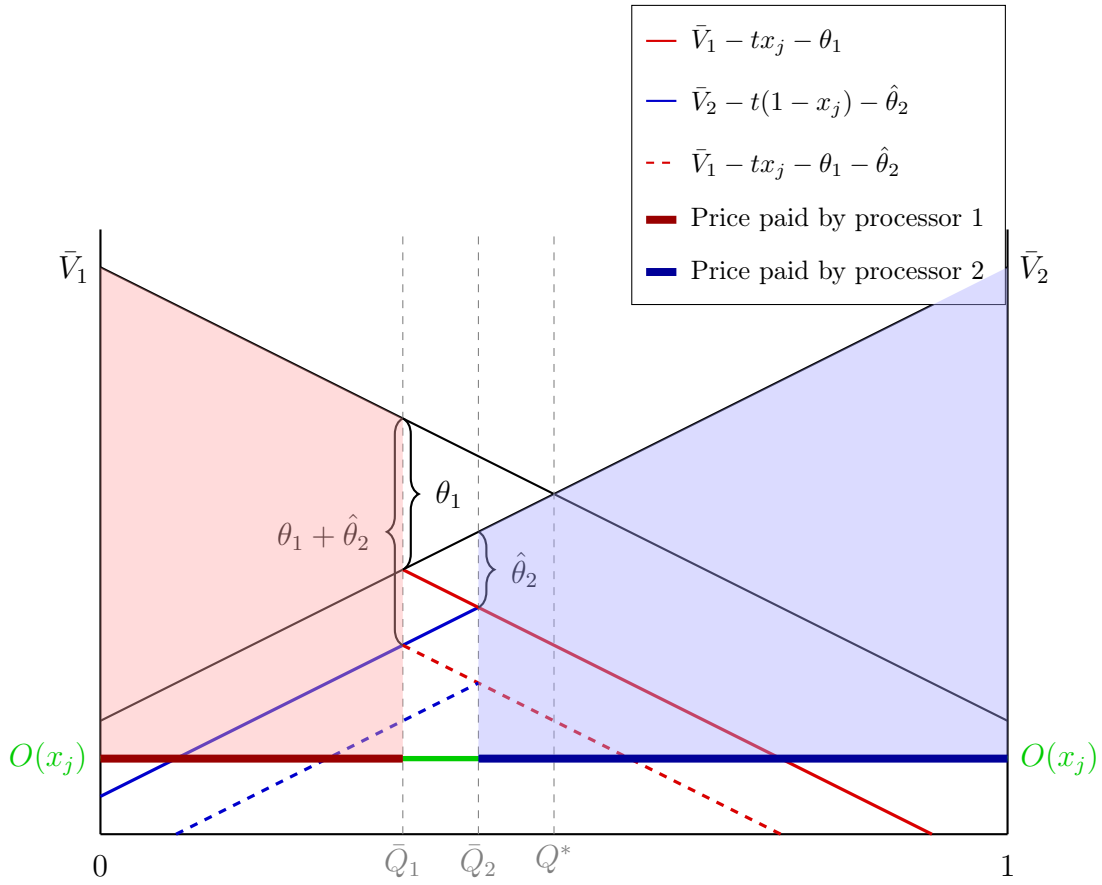
In the next subsections we illustrate two market conditions that prevent the full collapse of prices to sellers' outside option: excess processor capacity and heterogeneous seller outside options.

5.2 Excess Buyer Capacity Partially Restores Competition

Competition matters more when total processor capacity exceeds total feedlot production, which is the first failure mode of the CR condition identified after Proposition 2. Figure 2

[16] Lots in the interval $(\bar{Q}_1, \bar{Q}_2)$ are not purchased by either processor in this setting and exercise the outside selling option. In the spatial-competition literature, this is the configuration in which the local market is not fully covered, and each buyer operates as a local monopsonist within its own catchment area (Löfgren, 1986).

Figure 1: Bidding strategies for capacity-constrained processors in the Hotelling model with a uniform outside option



Notes: The figure shows how processor bids are attenuated for lots where the processor is not the high-valuation bidder due to profit guarantees on lots where it has the high valuation, given binding capacity constraints. Equilibrium contract prices are reduced to the sellers' outside option $O(x_j)$. Thick red and blue segments show the prices paid by processor 1 and processor 2 for the lots each buys; lots between $\bar{Q}_1$ and $\bar{Q}_2$ are sold at the outside option. Shaded areas represent processor profits.

retains the valuations of Figure 1 and the uniform outside option, but now capacities sum to more than available supply. Processor 1's smallest margin is $\lambda_1 = \theta_1 = V_1(\bar{Q}_1) - V_2(\bar{Q}_1)$, so it will not bid above the red line $V_1(x_j) - \theta_1$ to the right of $\bar{Q}_1$. Processor 2, despite a capacity of $1 - \bar{Q}_2$, can fill only the segment on $(\bar{Q}_1, 1]$ where its valuation exceeds processor 1's bid; its capacity does not bind, so its smallest margin is $\lambda_2 = \hat{\theta}_2 = 0$. Exactly as the reservation-demand logic predicts for a processor with slack capacity, it is willing to bid its full valuation on all lots in $[0, \bar{Q}_1]$.

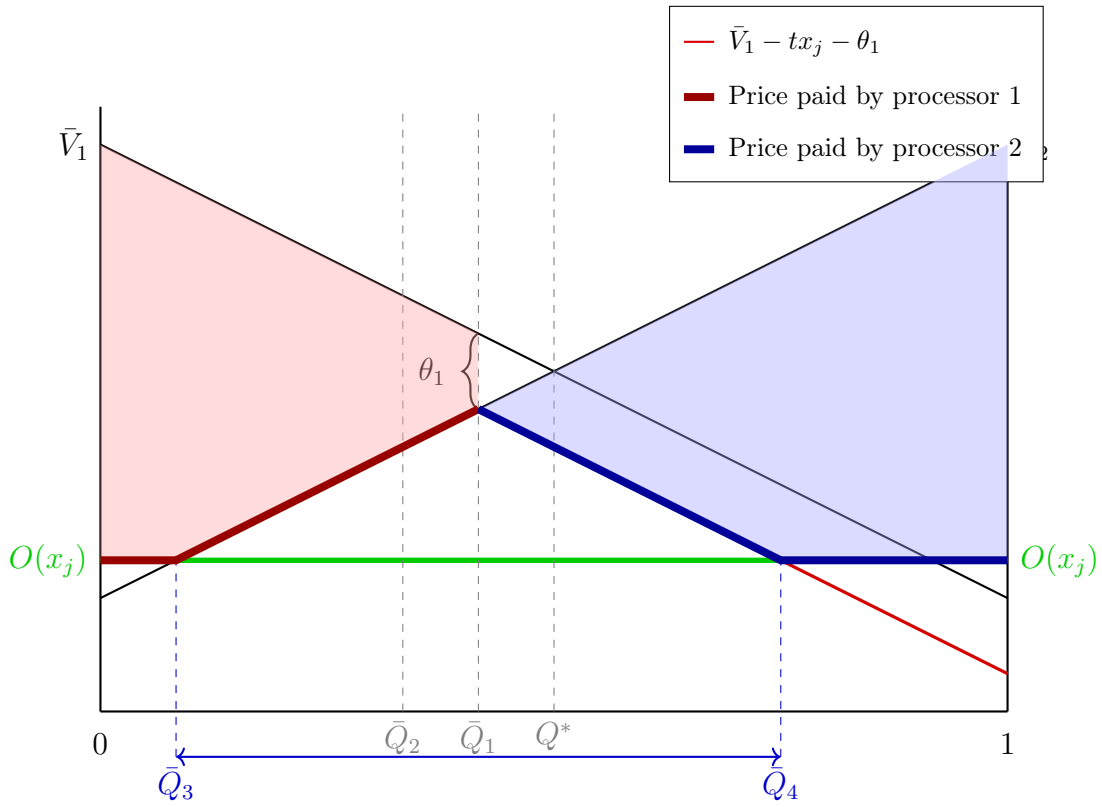
Feedlots on the interval $[\bar{Q}_3, \bar{Q}_4]$, where the rival processor's bid exceeds the outside option, receive prices above the outside option, while feedlots elsewhere receive $O(x_j)$. The shaded areas represent the processors' profits. Figure 2 thus illustrates the comparative statics result stated in Proposition 3: introducing excess capacity weakly raises prices for all sellers. Notably however, excess processor capacity does not produce the English second-price auction outcome. Although some sellers receive a price equivalent to the second-highest valuation, namely those in the interval $[\bar{Q}_3, \bar{Q}_1]$, the rest receive a lower price that may be bid down to their outside option.

5.3 Heterogeneous Outside Options Partially Restores Competition

Uniformity of the outside option also matters for the full collapse of prices to sellers' outside options because it is what guarantees that the CR condition holds. In this subsection, we introduce heterogeneous seller outside options and show how this condition prevents full collapse of prices due to the CR condition failing. Figure 3 presents a case where capacities sum to less than total supply, so there is no excess capacity. However, the outside option $O(x_j)$ (green line) declines from left to right and lies below both processors' valuations everywhere. It reflects a setting where feedlots near processor 1's plant have access to an attractive alternative outlet, such as a buyer in the adjacent region beyond the left market boundary. Perhaps due to more difficult transportation in the direction left of the market boundary, the outside option falls more steeply than either valuation line so that the margin over the outside option, $V_i(x_j) - O(x_j)$, rises with x_j for *both* processors. Valued at outside options, the best procurement opportunities of both processors therefore lie at the right end of the line. Their reservation demands overlap, and the CR condition fails.

The equilibrium allocation reverses the geography of Figure 1. Processor 2 fills its capacity with the feedlots nearest its plant on $[\bar{Q}_2, 1]$. Processor 1 fills its capacity $\bar{Q}_1$ with those feedlots in the middle portion, $[\bar{Q}_2 - \bar{Q}_1, \bar{Q}_2]$, of the market area. The feedlots on $[0, \bar{Q}_2 - \bar{Q}_1]$, nearest to processor 1's plant and holding the best outside options, sell to the outside market. Processor 1's smallest margin is λ_1, its margin on its marginal lot at

Figure 2: Bidding strategies in the Hotelling model with excess capacity

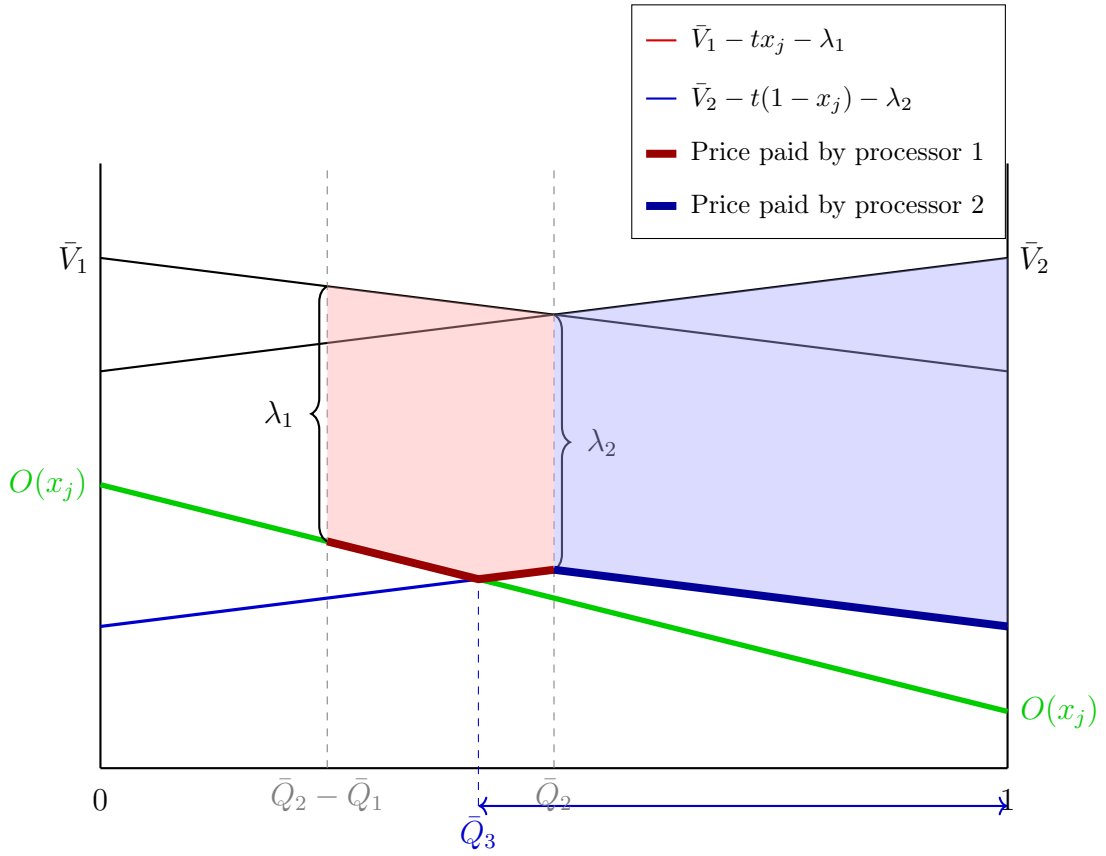


Notes: processor 2's capacity does not bind in this setting, so $\lambda_2 = 0$ is zero, and it bids up to its full valuation on lots in the interval $[0, \bar{Q}_1]$. Feedlots on the contested interval $[\bar{Q}_3, \bar{Q}_4]$ receive prices above the outside option. Thick red and blue segments show the prices paid by processor 1 and processor 2. Shaded areas are the processors' profits.

$\bar{Q}_2 - \bar{Q}_1$, which caps its bids on processor 2's segment at $V_1(x_j) - \lambda_1$ (red line). processor 2's smallest margin is λ_2, its margin on its marginal lot at $\bar{Q}_2$, which caps its bids on processor 1's segment at $V_2(x_j) - \lambda_2$ (blue line). We here have $\lambda_1 = \lambda_2$.

To the left of $\bar{Q}_3$ both bid lines lie below the outside option, so feedlots there are paid exactly $O(x_j)$ and the collapse of price received is complete. To the right of $\bar{Q}_3$ the rival's bid line rises above the outside option and the winning buyer must match it. Feedlots on $[\bar{Q}_3, 1]$ therefore receive prices strictly above their outside options, and the premium above outside option increases moving toward $x_j = 1$. The shaded areas are the processors' profits, with lower boundaries determined by the equilibrium contract prices.

Figure 3: Bidding strategies in the Hotelling model with declining outside options



Notes: The outside option declines with location more steeply than transport costs, so both processors' margins over the outside option rise toward $x_j = 1$ and the CR condition fails. Feedlots on $[0, \bar{Q}_2 - \bar{Q}_1]$ sell to the outside market. Feedlots on $[\bar{Q}_3, 1]$ receive price premia above their outside options. Thick red and blue segments show the prices paid by processor 1 and processor 2. Shaded areas are the processors' profits.

Again, the collapse of prices is partial, but the outcome is still the buyer-optimal equilibrium of Proposition 1: every price is the lowest at which no rival processor demands

the lot, and each processor earns at least λ_i on every lot it buys. The main conclusion remains: non-collusive price competition does little to support farm prices when overall production exceeds total capacity of the processors.

6 Discussion and Conclusion

This paper has shown how binding processor capacity constraints and heterogeneous buyer valuations of the farm product can interact to severely attenuate contract bidding in concentrated procurement markets and shift market surplus from farmers to their downstream buyers. The matching-theoretic formulation set forth in Sections 3 and 4 makes the argument precise: contract bidding with buyer offers selects the buyer-optimal Nash equilibrium (Proposition 1), and prices collapse fully to sellers' outside options whenever processors' capacity-constrained reservation demands do not compete for the same farm production (Proposition 2).

Although plant capacities are fixed at any point in time, they can be set strategically over longer time horizons by closing plants, such as occurred in the cattle industry in 2026 when Tyson Foods, Cargill, and JBS all announced plant closures or capacity reductions, following a period of steadily rising live cattle prices (Hanrahan, 2026). More subtle means also exist to credibly limit processing plants' capacity, including restricting plants' hours of operations, limiting the speeds at which plants operate, and so forth.[17] Proposition 3 gives this incentive a formal basis: any capacity reduction weakly lowers every seller's contract price, shifting the entire range of equilibrium prices against sellers. Given that binding capacity constraints direct a greater share of profits to processors, firms have clear short-run incentives to engage in such actions.

Given that the market conditions studied here—procurement markets featuring high buyer concentration, contract procurement through buyer offers, heterogeneous buyer valuation of farm production, and fixed plant capacity constraints that may bind—characterize many modern agricultural markets worldwide, the key question is what can be done to improve farmers' position in the market in these settings? A first answer is to recognize that results reported here were based on processors' short-run maximization of profits. No consideration was given to the long-run implications of depressing farm prices below their market valuation. As Crespi, Saitone, and Sexton (2012) have argued, such a strategy is likely not in buyers' long-run interests if the exercise of market power leads to the exit of farms that supply the essential input to processors' operations.

These authors showed that under certain market conditions, it was in buyers' interest to offer farm prices that were high enough to preserve the viability of the "stock" of farming operations, creating what they called a "symbiotic relationship" between buyers and sellers. In the language of our model, pursuit of short-run profits could turn a situation in which the

[17] In cattle processing, the speed at which a plant operates is known as "chain speed." Reducing the chain speed effectively reduces the plant's capacity to process cattle over a given time period.

CR condition holds—binding capacities and full price collapse—into one with excess capacity and higher contract prices due to the exit of farms. However, Crespi, Saitone, and Sexton (2012) also note that the conditions for a symbiotic buyer-seller relationship to emerge are not trivial. Buyers must have a long-run horizon and not discount future returns too much, and the market must not be prone to seller moral hazard, such as a seller defecting to a competing processor or an alternative crop despite a processor having paid the seller a price above its outside option over prior periods.

The analysis also motivates other strategies and policies that can support higher farm prices. Per proposition 3, any expansion of processing capacity within a market area weakly raises farm prices. Strategies such as the Ranchers First Initiatives proposed in the U.S. by the Trump Administration, which included loans to support regional processing and establishing processing cooperatives, in principle meet this objective.[18] Similar initiatives to expand beef processing in the U.S. were pursued by the Biden Administration.

An additional factor depressing farm prices in our model is that, consistent with most real-world contract bidding, processors make take-it-or-leave-it offers to farmers, which farmers then accept or reject. Allowing "negotiation" in this setting would not be helpful because processors would rationally reject any attempt by farmers to counteroffer with a higher price. The mechanism that would "turn the tables" in farmers' favor would be if farmers were able to post binding prices at which they were willing sell. Creating trading platforms that embodied this feature would face many practical difficulties, including matching buyers with sellers to maximize the market surplus, something the buyer-offer framework we study accomplishes, albeit while directing most or all of the surplus to the buyer.

References

Brandenburger, A., A. Friedenberg, and H.J. Keisler. 2008. "Admissibility in Games." *Econometrica* 76:307–352.

Butcher, E., and L. Schulz. 2025. "Production Dynamics and Disruption Responses in the Pork Supply Chain: A Structural Model of Hog and Pig Markets." *Journal of Agricultural and Applied Economics* 57:633–655.

Crawford, V.P., and E.M. Knoer. 1981. "Job Matching with Heterogeneous Firms and Workers." *Econometrica* 49:437–450.

[18]Increments to processing capacity would need to be efficient in the sense of having valuations on par with incumbent processors to play an effective pro-competitive role.

Crespi, J., and J. MacDonald. 2022. "Concentration in Food and Agricultural Markets." In *Handbook of Agricultural Economics*. Elsevier, vol. 6, pp. 4781–4843.

Crespi, J.M., T.L. Saitone, and R.J. Sexton. 2012. "Competition in U.S. Farm Product Markets: Do Long-Run Incentives Trump Short-Run Market Power?" *Applied Economic Perspectives and Policy* 34:669–695.

Deconinck, K. 2021. "Concentration and Market Power in the Food Chain." OECD Food, Agriculture and Fisheries Papers No. 151, OECD Publishing, Paris.

Dries, L., T. Reardon, and J.F.M. Swinnen. 2004. "The Rapid Rise of Supermarkets in Central and Eastern Europe: Implications for the Agrifood Sector and Rural Development." *Development Policy Review* 22:525–556.

Evans, A., and J. Sesmero. 2024. "Joint Choice of Location and Spatial Pricing Policy in a Mixed Market." https://ageconsearch.umn.edu/record/343956.

Faminow, M.D., and B.L. Benson. 1990. "Integration of Spatial Markets." *American Journal of Agricultural Economics* 72:49–62.

Fousekis, P. 2011. "Free-on-Board and Uniform Delivery Pricing Strategies in a Mixed Duopsony." *European Review of Agricultural Economics* 38:119–139.

Goodhue, R.E. 2011. "Food Quality: The Design of Incentive Contracts." *Annual Review of Resource Economics* 3:119–140.

—. 2027. "Price Setting in Agricultural Contracts." *Annual Review of Resource Economics* 19, In press.

Graubner, M., A. Balmann, and R.J. Sexton. 2011. "Spatial Price Discrimination in Agricultural Product Procurement Markets: A Computational Economics Approach." *American Journal of Agricultural Economics* 93:949–967.

Graubner, M., K. Salhofer, and C. Tribl. 2021. "A Line in Space: Pricing, Location, and Market Power in Agricultural Product Markets." *Annual Review of Resource Economics* 13:85–107.

Hamilton, S.F., and D.L. Sunding. 2021. "Joint Oligopsony-Oligopoly Power in Food Processing Industries: Application to the Us Broiler Industry." *American Journal of Agricultural Economics* 103:1398–1413.

Hanrahan, R. 2026. "Lawmakers Target Meatpacker Consolidation Amid High Prices." Farm Policy News, May 27.

Hatfield, J.W., and P.R. Milgrom. 2005. "Matching with Contracts." *American Economic Review* 95:913–935.

Hotelling, H. 1929. "Stability in Competition." *The Economic Journal* 39:41–57.

Hueth, B., and E. Ligon. 1999. "Producer Price Risk and Quality Measurement." *American Journal of Agricultural Economics* 81:512–524.

Jung, J., J. Sesmero, and R. Siebert. 2022. "A Structural Estimation of Spatial Differentiation and Market Power in Input Procurement." *American Journal of Agricultural Economics* 104:613–644.

Jung, J., J.P. Sesmero, and R. Siebert. 2026. "Using Pre- and Post-Entry Data to Identify the Effect of Ethanol Plant Entry on the Spatial Pattern of Corn Prices: A Study in Indiana." *American Journal of Agricultural Economics*, Sep., pp. .

Kelso, A.S., and V.P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* 50:1483–1504.

Kreps, D.M., and J.A. Scheinkman. 1983. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell Journal of Economics* 14:326–337.

Löfgren, K. 1986. "The Spatial Monopsony: A Theoretical Analysis." *Journal of Regional Science* 26:707–730.

MacDonald, J. 2024. "Concentration in U.S. meatpacking industry and how it affects competition and cattle prices." Amber Waves, U.S. Department of Agriculture, Economic Research Service.

MacDonald, J.M., and P. Korb. 2011. "Agricultural Contracting Update: Contracts in 2008." U.S. Department of Agriculture, Economic Research Service, Economic Information Bulletin No. 72.

Marx, L.M., and J.M. Swinkels. 1997. "Order Independence for Iterated Weak Dominance." *Games and Economic Behavior* 18:219–245.

Meemken, E.M., and M.F. Bellemare. 2020. "Smallholder Farmers and Contract Farming in Developing Countries." *Proceedings of the National Academy of Sciences* 117:259–264.

Mérel, P.R., R.J. Sexton, and A. Suzuki. 2009. "Optimal Investment in Transportation Infrastructure When Middlemen Have Market Power: A Developing-Country Analysis." *American Journal of Agricultural Economics* 91:462–476.

Moschini, G., and T.J. Smith. 2026. "Spatial Price Competition and Buyer Power in the U.S. Beef Packing Industry." *American Journal of Agricultural Economics* 108:771–798.

Moulin, H. 1979. "Dominance Solvable Voting Schemes." *Econometrica* 47:1337–1351.

Muth, M., G. Brester, J. Del Roccili, S. Koontz, B. Martin, N. Piggott, J. Taylor, T. Vukina, and M. Wohlgenant. 2005. "Spot and Alternative Marketing Arrangements in the Livestock and Meat Industries: Interim Report." U.S. Department of Agriculture, Grain Inspection, Packers and Stockyards Administration.

Osborne, M.J., and C. Pitchik. 1986. "Price Competition in a Capacity-Constrained Duopoly." *Journal of Economic Theory* 38:238–260.

Rogers, R.T., and R.J. Sexton. 1994. "Assessing the Importance of Oligopsony Power in Agricultural Markets." *American Journal of Agricultural Economics* 76:1143–1150.

Roth, A.E., and M.A.O. Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monographs, Cambridge: Cambridge University Press.

Saitone, T., K. Schaefer, and D. Scheitrum. 2021. "COVID-19 morbidity and mortality in U.S. meatpacking counties." *Food Policy* 101:102072.

Saitone, T.L., and R.J. Sexton. 2017. "Agri-Food Supply Chain: Evolution and Performance with Conflicting Consumer and Societal Demands." *European Review of Agricultural Economics* 44:634–657.

Sesmero, J., J. Balagtas, and M. Pratt. 2015. "The Economics of Spatial Competition for Corn Stover." *Journal of Agricultural and Resource Economics* 40:425–44.

Sexton, R.J., and T. Xia. 2018. "Increasing Concentration in the Agricultural Supply Chain: Implications for Market Power and Sector Performance." *Annual Review of Resource Economics* 10:229–251.

Shapley, L.S., and M. Shubik. 1971. "The Assignment Game I: The Core." *International Journal of Game Theory* 1:111–130.

Sheldon, I.M. 2017. "The Competitiveness Of Agricultural Product And Input Markets: A Review And Synthesis Of Recent Research." *Journal of Agricultural and Applied Economics* 49:1–44.

Squire, M. 2026. "Tyson Foods Beef Shake-Up Sparks Concern Across Cattle Country." *Successful Farming* August 14.

U.S. Government Accountability Office. 2009. "U.S. Agriculture: Retail Food Prices Grew Faster Than the Prices Farmers Received for Agricultural Commodities, but Economic Research Has Not Established That Concentration Has Affected These Trends." GAO Report No. GAO-09-746R, U.S. Government Accountability Office, Washington, D.C.

Wang, H.H., Y. Wang, and M.S. Delgado. 2014. "The Transition to Modern Agriculture: Contract Farming in Developing Economies." *American Journal of Agricultural Economics* 96:1257–1271.

Zhang, M., and R.J. Sexton. 2001. "FOB or Uniform Delivered Prices: Strategic Choice and Welfare Effects." *The Journal of Industrial Economics* 49:197–221.

Appendix A Non-Collusive Equilibrium

This appendix proves Propositions 1–3. Our approach follows the two-sided matching literature (Shapley and Shubik, 1971; Crawford and Knoer, 1981; Kelso and Crawford, 1982; Hatfield and Milgrom, 2005). In the contract market, processors' valuations of lots are additive within capacity, and feedlots care only about the price they receive. We show in this first appendix that non-collusive equilibria, the outcomes of Nash equilibria, are exactly the outcomes satisfying conditions (i)–(iii) of Section 3.2, and that the buyer-optimal outcome is well defined. We then introduce a processor-proposing bidding process that converges to the buyer-optimal equilibrium (Appendix B). Appendix C finishes the proofs. As in the main text, we often refer to buyers as processors and sellers as feedlots, while noting that the results apply to contract procurement of agricultural products more generally.

A.1 Demand and the no-ties assumption

There are n processors and m feedlots. Processor i can process at most q_i lots and values lot j at $V_{ij} > 0$ Feedlot j has outside option $O_j > 0$. Given a price vector $\boldsymbol{p} = (p_1, \ldots, p_m)$, processor i's profit from a set C_i of at most q_i lots is $\pi^i(C_i; \boldsymbol{p}) = \sum_{j \in C_i}(V_{ij} - p_j)$, as in Equation (1). A set C_i^* with $|C_i^*| \leq q_i$ is a *most profitable set* for processor i at $\boldsymbol{p}$ if $\pi^i(C_i^*; \boldsymbol{p}) \geq \pi^i(C_i; \boldsymbol{p})$ for every C_i with $|C_i| \leq q_i$.

Let $\mu_{ij} = V_{ij} - p_j$ represent processor i's margin on lot j. Because profit is additive, a set C_i^* with $|C_i^*| \leq q_i$ is a most profitable set for processor i at $\boldsymbol{p}$ if and only if:

(a) $\mu_{ij} \geq 0$ for every $j \in C_i^*$ (no lot in C_i^* is worth dropping);
(b) $\mu_{ij} \geq \mu_{ik}$ for every $j \in C_i^*$ and $k \notin C_i^*$ (no lot in C_i^* is worth swapping out);
(c) if $|C_i^*| < q_i$, then $\mu_{ik} \leq 0$ for every $k \notin C_i^*$ (no lot is worth adding).

These are the three single-lot changes ruled out by condition (ii) in Section 3.2. All of them are weak linear inequalities in $\boldsymbol{p}$, which the following arguments use.

An *allocation* $C = (C_1, \ldots, C_n)$ gives disjoint sets of lots to the processors with $|C_i| \leq q_i$, and lots outside $\cup_i C_i$ are sold to the outside option. The economic surplus associated with allocation C is

$$W(C) \;=\; \sum_{i=1}^{n} \sum_{j \in C_i} V_{ij} \;+\; \sum_{j \notin \cup_i C_i} O_j \;. \tag{A.1}$$

Assumption A.1 (No ties)**.** Distinct allocations generate distinct surpluses: $W(C) \neq W(C')$ whenever $C \neq C'$.

This is the formal version of the no-ties condition in Section 3.1. When $\boldsymbol{p} = \boldsymbol{O}$, a tie

means that multiple allocations can lead to the same value of surplus, and this assumption rules it out. It guarantees that each processor has a unique most profitable set at $\boldsymbol{O}$ (its reservation demand S_i) and that the surplus-maximizing (i.e., efficient) allocation is unique. Ties at equilibrium prices, by contrast, are created by competition and are not ruled out.

A.2 Nash equilibrium and non-collusive equilibrium

Definition A.1 (Non-collusive equilibrium)**.** A *non-collusive equilibrium* $(C^*, \boldsymbol{p})$ is the outcome of a pure-strategy Nash equilibrium of the bidding game, where C^* is the allocation of lots, and $\boldsymbol{p}$ records the accepted bid on every lot sold to a processor and the outside option O_j on every lot sold at the outside option.

In a Nash equilibrium, each processor sets its own bids to maximize its profit, taking its rivals' bids as given. The next lemma expresses the equilibrium in terms of three conditions on prices.

Lemma A.1 (Characterization of non-collusive equilibrium)**.** *An allocation C^* and a price vector $\boldsymbol{p}$ form a non-collusive equilibrium if and only if*

(i) $p_j \geq O_j$ for every lot j sold to a processor;
(ii) C_i^* is a most profitable set for every processor i at $\boldsymbol{p}$;
(iii) $p_j = O_j$ for every lot j sold at the outside option.

Proof. ($\Rightarrow$) Consider a Nash equilibrium with outcome $(C^*, \boldsymbol{p})$. Conditions (i) and (iii) hold because a feedlot accepts only bids at or above its outside option. Every bid on lot j is at most p_j, since the feedlot accepts it as the best offer. If condition (ii) failed for processor i, it could earn $\pi^i(S; \boldsymbol{p}) > \pi^i(C_i^*; \boldsymbol{p})$ with some alternative set of acquisitions S. Processor i could then bid $b_{ij} = p_j + \varepsilon$ on every lot j in S and bid 0 on all other available lots, win exactly the lots in S, and earn $\pi^i(S; \boldsymbol{p}) - \varepsilon|S| > \pi^i(C_i^*; \boldsymbol{p})$ for small ε, which contradicts that $(C^*, \boldsymbol{p})$ represents the non-collusive outcome of a Nash equilibrium.

($\Leftarrow$) Take $(C^*, \boldsymbol{p})$ satisfying (i)–(iii). Let every processor, including those that do not buy lot j, bid p_j on every lot sold to a processor and nothing on the other lots, and let each such lot go to its processor in C^*. Each processor i wins C_i^* and earns $\pi^i(C_i^*; \boldsymbol{p})$. In any deviation, processor i pays at least p_j for every lot j it wins as a result of outbidding its rivals or the outside option. Lots beyond capacity add nothing, so its profit is at most $\max_{|S| \leq q_i} \pi^i(S; \boldsymbol{p}) = \pi^i(C_i^*; \boldsymbol{p})$ by (ii). No deviation is profitable, so these bids form a Nash equilibrium with outcome $(C^*, \boldsymbol{p})$. □

A.3 The buyer-optimal equilibrium

Lemma A.2 (Buyer-optimal equilibrium)**.** *Let Assumption A.1 hold and suppose a non-collusive equilibrium exists. Then:*

(a) every non-collusive equilibrium contains the efficient allocation C^*;

(b) the set P of equilibrium price vectors has a least element $\boldsymbol{p}^*$ such that $\boldsymbol{p}^* \leq \boldsymbol{p}$ coordinatewise for every $\boldsymbol{p} \in P$;

(c) at $(C^*, \boldsymbol{p}^*)$ every processor earns at least as much, and every feedlot receives at most as much, as at any other non-collusive equilibrium.

Proof. (a) Let $(C, \boldsymbol{p})$ be a non-collusive equilibrium and C' any allocation. Write $A = \cup_i C_i$ and $A' = \cup_i C'_i$ for the sets of lots sold to processors, and $A \setminus A'$ for the lots in A but not in A' and $A' \setminus A$ for the lots in A' but not in A. Substituting $V_{ij} = \mu_{ij} + p_j$ into (A.1) yields

$$W(C) - W(C') \;=\; \sum_i \big[\pi^i(C_i; \boldsymbol{p}) - \pi^i(C'_i; \boldsymbol{p})\big] \;+\; \sum_{j \in A \setminus A'} (p_j - O_j) \;-\; \sum_{j \in A' \setminus A} (p_j - O_j)\ . \quad \text{(A.2)}$$

By condition (ii) of Lemma A.1, C_i is a most profitable set for every processor i at $\boldsymbol{p}$, so the first term in Equation (A.2) is nonnegative. By condition (i) of Lemma A.1, $p_j \geq O_j$ for every lot j sold to a processor, so the second term in Equation (A.2) is nonnegative. By condition (iii) of Lemma A.1, lots not in $A = \cup_i C_i$ get sold at the outside option, and $p_j = O_j$ for every such lot j, so the third term equals zero. Therefore, we have assignment C maximizes W, and by Assumption A.1 it is C^*.

(b) We need one price vector in P that is lowest in every element p_j. Every $\boldsymbol{p} \in P$ satisfies $O_j \leq p_j \leq \max_i V_{ij}$, so P is bounded. By part (a), P is the set of price vectors satisfying conditions (i)–(iii) of Lemma A.1 at C^*, and each of these is a weak inequality in $\boldsymbol{p}$ (condition (ii) through the demand conditions (a)–(c)). A bounded set defined by weak inequalities contains its boundary, so there exists some $\boldsymbol{p}^* \in P$ that minimizes the total price $\sum_j p_j$ over P.

For any $\boldsymbol{p}, \boldsymbol{p}' \in P$, we write their coordinatewise minimum as $\hat{\boldsymbol{p}} = \boldsymbol{p} \wedge \boldsymbol{p}'$, that is, $\hat{p}_j = \min\{p_j, p'_j\}$ for every lot j. We show that $\hat{\boldsymbol{p}}$ is also an equilibrium price vector. At $\hat{\boldsymbol{p}}$, every margin takes its larger value under $\boldsymbol{p}$ and $\boldsymbol{p}'$, $\mu_{ij}(\hat{\boldsymbol{p}}) = \max\{\mu_{ij}(\boldsymbol{p}), \mu_{ij}(\boldsymbol{p}')\}$. Processor i holds the same set C_i^* under $\boldsymbol{p}$ and $\boldsymbol{p}'$, so the demand conditions (a)–(c) hold under both vectors, and taking the larger of two margins preserves all of them,[19] so C_i^* is a most profitable set for every processor i at $\hat{\boldsymbol{p}}$. We also have $p_j \geq O_j$ for every lot j sold to a processor and $p_j = O_j$ for every lot j sold at the outside option, so by Lemma A.1, $(C^*, \hat{\boldsymbol{p}})$ is a non-collusive

[19]Specifically, (a) If $\mu_{ij} \geq 0$ under both vectors, the larger is nonnegative. (b) If $\mu_{ij} \geq \mu_{ik}$ under both, then $\max\{\mu_{ij}(\boldsymbol{p}), \mu_{ij}(\boldsymbol{p}')\} \geq \max\{\mu_{ik}(\boldsymbol{p}), \mu_{ik}(\boldsymbol{p}')\}$. (c) If $\mu_{ik} \leq 0$ under both, the larger is nonpositive.

equilibrium.

To see why $\boldsymbol{p}^*$ is lowest in every element p_j, suppose some $\boldsymbol{p} \in P$ has $p_j < p_j^*$ for at least one lot j. Then $\hat{\boldsymbol{p}} = \boldsymbol{p}^* \wedge \boldsymbol{p}$ lies in P, satisfies $\hat{\boldsymbol{p}} \leq \boldsymbol{p}^*$, and is strictly lower in lot j, so $\sum_j \hat{p}_j < \sum_j p_j^*$. This contradicts that $\boldsymbol{p}^*$ minimizes the total price over P. So $\boldsymbol{p}^*$ is the vector we want.

(c) All equilibria share the allocation C^*. As prices rise, each processor's profit falls, and each feedlot's payoff rises. □

As in Shapley and Shubik (1971), equilibrium prices thus form a range whose bottom is the buyer-optimal equilibrium. Lemma A.2 shows only that such an equilibrium is well defined. The next appendix proves that an equilibrium exists and shows how the market reaches the buyer-optimal one.

Appendix B The Processor-Proposing Bidding Process

We adapt the salary-adjustment process of Crawford and Knoer (1981) and KC (Kelso and Crawford, 1982) using the matching-with-contracts notation of Hatfield and Milgrom (2005). The process entails the logic of an ascending sequence of processor bids and seller rejections that starts from the feedlots' outside options. Though characterized as a sequential bidding process, it is only a mechanism to compute where processor bidding ends. Under complete information, each processor knows how its competitors will behave and computes the endpoint of the process without submitting those bids in reality. Processors submit their bids only once, informed by the simulated bidding process.

B.1 The process

We have prices move on a grid and fix an increment $\delta > 0$ as the unit of measurement. Lot j can thus be priced at $O_j, O_j + \delta, O_j + 2\delta, \ldots$, up to $\max_{i,j} V_{ij} + \delta$.[20] A *bid* is a triple (i, j, b): processor i offers price b for lot j. Let X be the finite set of all bids. Given a set of bids $Y \subseteq X$:

- processor i *chooses* its most profitable set by taking its cheapest bid in Y on each lot and keeping the (up to) q_i lots with the highest positive margins at those prices;
- feedlot j *chooses* the highest bid on lot j in Y (ties broken by processor index).

[20] We choose δ so that no processor is indifferent between two sets of lots at grid prices. Profit differences at grid prices take the form $\Delta - k\delta$ with $\Delta \neq 0$ by Assumption A.1, where k is an integer, so this excludes only countably many values of δ.

The bids a buyer or seller does not choose are *rejected*, and we let $R_P(Y)$ and $R_F(Y)$ collect the bids rejected by processors and by feedlots.

Lemma B.1 (Monotone rejections)**.** *If $Y \subseteq Y'$, then $R_F(Y) \subseteq R_F(Y')$ and $R_P(Y) \subseteq R_P(Y')$.*

Proof. A feedlot's best bid can only improve when more bids arrive, so a bid it rejects stays rejected. For a processor, suppose it rejects (i, j, b) from Y. If Y' contains a cheaper bid on lot j, it rejects (i, j, b) again. Otherwise, its margin on lot j is unchanged, while the added bids in $Y' \setminus Y$ can only raise its margins on other lots, so lot j still does not enter its most profitable set. □

Iteration. Let Y_P^t be the bids feedlots have not rejected through round t, and Y_F^t the bids processors have offered through round t. Y_P^t starts at $Y_P^0 = X$. In round t, processors choose from Y_P^{t-1}, so $Y_F^t = X \setminus R_P(Y_P^{t-1})$, and feedlots then choose from Y_F^t and update $Y_P^t = X \setminus R_F(Y_F^t)$. By Lemma B.1, Y_F^t grows and Y_P^t shrinks. Each time feedlot j rejects processor i, processor i's lowest permitted price on lot j, which we call ρ_{ij}, rises by δ. Three properties follow directly from the choice rules and Lemma B.1. First, a bid a feedlot holds is offered again in the next round, so the price a feedlot holds never falls. Second, processor i is rejected on lot j only when feedlot j holds a rival bid at least as high. Third, processors never offer bids with negative margin, so $\rho_{ij} \leq V_{ij} + \delta$, and the process stops after finitely many rounds. When it stops, each feedlot sells to the processor whose bid it holds at that price, and a feedlot holding no bid sells outside. Denote this *final outcome* by $(C^\delta, \boldsymbol{p}^\delta)$, where C_i^δ is processor i's most profitable set at its permitted prices $\boldsymbol{\rho}_i$.[21]

B.2 Existence and selection

Lemma B.2 (Existence)**.** *As $\delta \to 0$, the final outcomes converge to an outcome satisfying conditions (i)–(iii) of Lemma A.1, so a non-collusive equilibrium exists, and a pure-strategy Nash equilibrium exists.*

Proof. The final outcome satisfies (i) and (iii) of Lemma A.1 by construction. For (ii), processor i's set C_i^δ is a most profitable set at its permitted prices $\boldsymbol{\rho}_i$, which equal $\boldsymbol{p}^\delta$ on the lots it holds. On a lot it does not hold, either it never bid ($\rho_{ij} = O_j \leq p_j^\delta$) or it was last

[21] The two-lot market of footnote 13, with $O_1 = O_2 = 1$ and a small increment δ, illustrates the idea. In round 1, both processors bid 1 on lot 1; feedlot 1 holds one bid and rejects the other. Rejections on lot 1 alternate between the processors as the price rises. Once the price of lot 1 passes 6, processor 2's margin $9 - p_1$ falls below its margin on lot 2, which is $4 - 1 = 3$, so it switches to lot 2 at price 1, and no further rejection occurs. The process stops within one increment of the buyer-optimal prices $(6, 1)$.

rejected at $\rho_{ij} - \delta$ by a feedlot ($\rho_{ij} \le p_j^\delta + \delta$). Raising its prices on lots it does not hold to $\max\{\rho_{ij}, p_j^\delta\}$ only makes those lots less attractive, so C_i^δ is a most profitable set at a price vector within δ of $\boldsymbol{p}^\delta$. There are finitely many allocations and prices are bounded, so along a subsequence $\delta_k \to 0$ the allocation is constant, and prices converge. The conditions (i)–(iii) are weak inequalities, so they are preserved in the limit. □

Proposition B.1 (The process favors the processors)**.** *Let $(C^y, \boldsymbol{p}^y)$ be any non-collusive equilibrium whose prices lie on the grid. Throughout the process, $\rho_{ij} \le p_j^y$ for every processor i and every lot $j \in C_i^y$. Consequently, at the final outcome $(C^\delta, \boldsymbol{p}^\delta)$, every processor earns at least as much as at $(C^y, \boldsymbol{p}^y)$.*

Proof. We first prove that no feedlot ever rejects a bid (i, j, b) with $j \in C_i^y$ and $b \ge p_j^y$. Suppose otherwise, and take the first round in which this happens: feedlot j rejects (i, j, b) and holds a rival bid (k, j, b') with $b' \ge b \ge p_j^y$, while lot j is in C_i^y. Since this is the first such rejection, processor k can still bid p_l^y on each of its own lots $l \in C_k^y$. Processor k chose (k, j, b') from its available bids, so by Lemma B.1 it also chooses it from the smaller set $Z = \{(k, j, b')\} \cup \{(k, l, p_l^y) : l \in C_k^y\}$. Let $A \ni j$ be its chosen set from Z. Since $A \ne C_k^y$ and there is no tie at grid prices,

$$\pi^k(A; \boldsymbol{p}^y) \;\ge\; \pi^k\big(A;\, b', \boldsymbol{p}^y_{-j}\big) \;>\; \pi^k\big(C_k^y;\, b', \boldsymbol{p}^y_{-j}\big) \;=\; \pi^k\big(C_k^y; \boldsymbol{p}^y\big),$$

where the first inequality uses $b' \ge p_j^y$ and the equality uses $j \notin C_k^y$. Thus processor k strictly prefers A to C_k^y at the prices $\boldsymbol{p}^y$, contradicting condition (ii) that C_k^y is a most profitable set for processor k at $\boldsymbol{p}^y$.

Without such rejections, each processor i can always bid p_j^y on each $j \in C_i^y$, so $\rho_{ij} \le p_j^y$ there. At the final outcome, C_i^δ is a most profitable set at $\boldsymbol{\rho}_i$, so

$$\pi^i(C_i^\delta; \boldsymbol{p}^\delta) = \pi^i(C_i^\delta; \boldsymbol{\rho}_i) \ge \pi^i(C_i^y; \boldsymbol{\rho}_i) \ge \pi^i(C_i^y; \boldsymbol{p}^y)\ .$$

This proves the proposition. □

By Lemma B.2, as $\delta \to 0$, the final outcomes converge along a subsequence to a non-collusive equilibrium that carries the allocation C^* by Lemma A.2(a) and so has prices at least $\boldsymbol{p}^*$. By Proposition B.1, the final prices exceed $\boldsymbol{p}^*$ by at most a few increments, which also vanish as $\delta \to 0$. The limit of the process is therefore the buyer-optimal equilibrium $(C^*, \boldsymbol{p}^*)$, and Proposition B.1 carries over to the market with continuous prices.

Appendix C Proofs

By Lemma A.1, non-collusive equilibria, the outcomes of Nash equilibria, are exactly the outcomes satisfying conditions (i)–(iii), so we argue in terms of these conditions. Assumption A.1 holds throughout.

Proof of Proposition 1. (i) An equilibrium exists by Lemma B.2. (ii) By Lemma A.2, all equilibria share the allocation C^*, and $(C^*, \boldsymbol{p}^*)$ gives every processor its highest profit and every feedlot its lowest price. (iii) The processor-proposing process starts from the outside options, lets processors' bids drive prices up, and converges to $(C^*, \boldsymbol{p}^*)$ as $\delta \to 0$ (Proposition B.1 and the discussion following it). The proof of Lemma A.1 gives the Nash equilibrium bids that deliver this outcome. □

Proof of Proposition 2. Let S_i be processor i's reservation demand, its unique most profitable set at $\boldsymbol{O}$ by Assumption A.1.

($\Leftarrow$) Suppose the CR holds. The sets $S_1, \ldots, S_n$ are disjoint, and $(S_1, \ldots, S_n)$ with prices $\boldsymbol{O}$ satisfies Lemma A.1 (i)–(iii). (i) and (iii) hold trivially, and (ii) holds by the definition of S_i. No equilibrium price can be below $\boldsymbol{O}$, so this is the buyer-optimal equilibrium and every price equals the outside option.

($\Rightarrow$) We prove the contrapositive. If the CR fails, some lot j lies in $S_i \cap S_k$ with $i \neq k$. An equilibrium at prices $\boldsymbol{O}$ would require $C_i = S_i$ and $C_k = S_k$ by Lemma A.1(ii), allocating lot j to two processors. So no equilibrium has prices $\boldsymbol{O}$, and the buyer-optimal equilibrium pays some feedlot strictly more than its outside option. □

Proof of Proposition 3. We prove the result for a one-unit increase in one processor's capacity. Any expansion is a sequence of such one-unit increases, so the proof can be extended to any increase in capacity. A reduction is an expansion read backwards, so it weakly lowers every price.

Compare the market E with the market E' in which processor k's capacity rises from q_k to $q_k + 1$. Mark objects of E' with primes. From any set of bids, processor k's $q_k + 1$ best lots include its q_k best, so it rejects weakly fewer bids, $R'_P(Y) \subseteq R_P(Y)$ for every Y.

Run the processor-proposing bidding process separately in E and E' from the same starting point $Y_P^0 = Y_P'^0 = X$, where no bid has yet been rejected. We show by induction that $Y_F'^t \supseteq Y_F^t$ and $Y_P'^t \subseteq Y_P^t$ in every round. For $t = 1$, $Y_F'^1 = X \setminus R'_P(X) \supseteq X \setminus R_P(X) = Y_F^1$ because $R'_P \subseteq R_P$. Then $Y_P'^1 = X \setminus R_F(Y_F'^1) \subseteq X \setminus R_F(Y_F^1) = Y_P^1$ by Lemma B.1. For the inductive step, suppose $Y_P'^{t-1} \subseteq Y_P^{t-1}$. Then $R'_P(Y_P'^{t-1}) \subseteq R'_P(Y_P^{t-1}) \subseteq R_P(Y_P^{t-1})$, first by Lemma B.1 and then because $R'_P \subseteq R_P$. We thus have $Y_F'^t \supseteq Y_F^t$, and Lemma B.1 gives $Y_P'^t \subseteq Y_P^t$ as before.

When both processes have stopped, every bid offered in E has also been offered in E'. Since a feedlot's final price is the highest bid it has received, or O_j if none, we have $p_j'^{\delta} \geq p_j^{\delta}$ for every j. As $\delta \to 0$, both final outcomes converge to the buyer-optimal equilibria of their markets, and the inequality holds in the limit. $\square$