

Collusion in Brokered Markets

John William Hatfield
Scott Duke Kominers
Richard Lowery

Working Paper 20-023



Collusion in Brokered Markets

John William Hatfield

University of Texas at Austin

Scott Duke Kominers

Harvard Business School

Richard Lowery

University of Texas at Austin

Working Paper 20-023

Copyright © 2019, 2020 by John William Hatfield, Scott Duke Kominers, and Richard Lowery

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Funding for this research was provided in part by Harvard Business School, the National Science Foundation (grants CCF-1216095 and SES-1459912), the Harvard Milton Fund, and both the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications.

Collusion in Brokered Markets*

John William Hatfield
McCombs School of Business
University of Texas at Austin

Scott Duke Kominers
Harvard Business School &
Department of Economics
Harvard University

Richard Lowery
McCombs School of Business
University of Texas at Austin

July 8, 2020

Abstract

The U.S. residential real estate agency market presents a puzzle for economic theory: commissions on real estate transactions have remained high for decades even though entry is frequent and costs are low. We model the real estate agency market, and other brokered markets, as a game in which brokers first post prices for customers and then choose which other agents to work with. We show that prices appreciably higher than the competitive level can be sustained (for a fixed discount factor) regardless of the number of brokers by using strategies that condition working with a broker on that broker's posted prices.

JEL Classification: D43, L13, L4, R39

Keywords: Real estate, Repeated games, Collusion, Antitrust, Brokered markets

*The authors thank Panle Jia Barwick, Peter Blair, Andrew F. Daughety, Zoe DeStorres, Drew Fudenberg, Sonia Gilbukh, Ed Glaeser, Paul Goldsmith-Pinkham, Oliver Hart, Louis Kaplow, Charles Nathanson, Michael Ostrovsky, Jennifer F. Reinganum, Kathy E. Spier, Abraham Wickelgren, Maisy Wong, and several conference and seminar audiences for helpful comments. Kominers is grateful for the support of the National Science Foundation (grants CCF-1216095 and SES-1459912), the Harvard Milton Fund, and both the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications. Any comments or suggestions are welcome and may be emailed to john.hatfield@utexas.edu.

1 Introduction

The real estate industry in the United States is characterized by widespread price coordination: The brokerage fee is typically 6%, with half of the fee going to the buyer’s agent and half going to the seller’s agent. There is little evidence that the 6% fee represents the true cost of facilitating a real estate transaction: it has been constant for decades (despite substantial changes in the technology used), varies with neither market conditions nor the price of the house being sold, and is appreciably higher than in many other countries (Delcours and Miller, 2002). Indeed, a number of lawsuits have recently been filed arguing that realtors conspire to keep commissions high.¹ Yet the market for providing residential real estate brokerage services is both highly unconcentrated and quite easy to enter (Beck et al., 2012), and so we might naturally expect price competition to drive down fees;² indeed, as Hsieh and Moretti (2003) remark in their influential study of real estate brokerage in the U.S., “the apparent uniformity of commission rates presents an enormous puzzle”.³

We provide a potential explanation for how the real estate brokerage industry can maintain high prices even in the presence of many independent brokers.⁴ We model the market for brokerage services as a repeated extensive form game: In each period, a continuum of buyers and sellers seek to buy and sell houses; buyers and sellers, however, are unable to transact directly and must instead work through agents. Each agent offers a buyer price and a seller

¹See *Moehrl v. NAR et al.* (U.S. Dist. Ct. N.D. Ill., Case No. 1:19-cv-01610), *Sawbill Strategic, Inc. v. NAR et al.* (U.S. Dist. Ct. N.D. Ill., Case No. 1:19-cv-02544), and *Sitzer and Winger v. NAR et al.* (U.S. Dist. Ct. W.D. Mo., Case No. 4:19-cv-00332); the Department of Justice has also filed a statement of interest in the last case.

²Indeed, real estate agents themselves have noted that restricting price competition is necessary to maintain high prices: For instance, a number of realtors from Montgomery County, Maryland, held a dinner party at the Congressional Country Club in Bethesda and discussed raising real estate commissions to 7%; this behavior was eventually investigated and prosecuted by the Department of Justice (see *U.S. v. Foley et al.*, 598 F.2d 1323 (4th Cir. 1979)). The realtors understood that raising the commission rate to 7% would only be effective if all the realtors participated, with John T. Carruthers stating that “if we do not stay at seven percent, then it would be a slide back and... no one could get seven percent, because the competition would hurt us.”

Of course, we are far from the first to note that industry competitors may wish to work together to maintain high prices; see, e.g., Smith (1776):

People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices.

³Similarly, Levitt and Syverson (2008b) and Bernheim and Meer (2013) argue that real estate agents provide poor service at high prices despite effectively free entry into real estate agency. Additionally, Barwick and Pathak (2015) argue that the current market structure is inefficient, with excessive commissions and too many agents; see also work by Barwick et al. (2017).

⁴Technically, there is a distinction between real estate brokers and real estate agents; specifically, a real estate broker can have real estate agents as employees while real estate agents cannot. Real estate agents can, however, operate independently. Barriers to entry are somewhat higher for brokers, but agents are able to compete directly with groups of brokers and thus overall barriers to entry in this market are low.

price for intermediation services; buyers and sellers then choose agents having observed these menus of prices. Once buyers and sellers have agents, each agent decides which other agents he is willing to work with; facilitating a transaction between a buyer and a seller requires both the buyer’s agent and the seller’s agent to be willing to work with each other.

In our setting, agents can maintain high prices by refusing to work with any agent who undercuts either of the “agreed upon” buyer and seller prices. This endogenously lowers the quality of a price deviator: a price deviator can no longer facilitate transactions between his buyers and another agent’s sellers (or between his sellers and another agent’s buyers). This implies that cutting prices by a small amount is *not* enough to attract buyers and sellers, as they understand that any price deviator will find it much more difficult to facilitate transactions. As a result, it is possible to maintain prices above marginal cost while ensuring that a price deviator who does attract buyers and sellers will not profit from his actions.⁵

Of course, it must be incentive compatible for agents to refuse to work with a price deviator. Here, we rely on the repeated interactions among agents: A non-price deviating agent is willing to forgo working with a price deviator today if he is sufficiently rewarded for doing so in the future. To see why, we note first that the non-price deviating agent’s forgone profits are small, since the non-price deviating agent only has a small fraction of the buyers and sellers. Second, the non-price deviating agent is incentivized to follow the prescribed punishment strategy as, if he does not, future play reverts to a equilibrium in which he obtains no profits in every future period. By contrast, if every non-price deviating agent punishes the price deviator as prescribed, future play moves to a collusive punishment phase in which every agent other than the price deviator obtains positive rents. Thus, a non-price deviating agent will be willing to forgo working with a price deviator today for reasonably high discount factors—even as the number of agents becomes arbitrarily large.

One might be incredulous that real estate agents would, in fact, engage in such behavior. However, such behavior is well-documented: In one particularly harrowing tale, [Birger and Caplin \(2004\)](#) reported that the proprietress of a new discount realty, You Win Realty, was not only boycotted by other realtors in the area but also harassed by threatening phone calls and the like; indeed, one competing realtor threatened to report her to the Federal Trade Commission(!). Taking a broader view, a report by the [Department of Justice and Federal Trade Commission \(2007\)](#) documented that such “steering” behavior (away from discount realtors) is commonly reported and an earlier investigation by the [Federal Trade Commission \(1983\)](#) found that 84% of alternative brokers reported

⁵We do not specifically address here why a fixed commission rate of 6% is used. But simple pricing schemes such as this are instrumental in maintaining collusion in environments in which information is not perfect; see [Athey and Bagwell \(2001\)](#), [Athey et al. \(2004\)](#), and [Kang and Lowery \(2014\)](#).

frequent or occasional “refusals by other brokers to show homes listed by [their] business,” with 49% saying that this was a frequent issue. Relatedly, [Barwick et al. \(2017\)](#) documented that properties with lower commission rates are both less likely to sell and take longer to sell.⁶

Our equilibrium is necessarily more complex than one constructed with simple penal codes à la [Abreu \(1988\)](#). In repeated normal form games, [Abreu \(1988\)](#) demonstrated that simple penal codes are sufficient for implementing maximally collusive strategies. In repeated extensive-form game settings such as ours, by contrast, simple penal codes may not be sufficient for characterizing maximal collusion since in-period punishment may play a role in supporting subgame-perfect equilibria. As [Mailath et al. \(2017\)](#) noted, the analysis of repeated extensive form games may require more involved responses to deviations; the need to reward within-period punishments implies that rewarding agents in future periods may be more important than just punishing the initial deviator as much as possible.⁷ Thus, in our model, we construct an equilibrium which promises sufficient rewards in future periods to each non-price-deviating agent to incentivize him to punish a price-deviator in-period.

Our analysis not only explains how real estate agents may maintain prices appreciably above cost, but also enables us to assess different proposed policies. [Han and Hong \(2011\)](#) investigated “rebate bans,” laws that prohibit buyers’ agents from sharing their commissions with buyers (although agents are still allowed to pay buyers’ closing costs); ten states have rebate bans in effect. We show that rebate bans facilitate higher commissions even though they still offer room for real estate agents to reduce buyers’ (closing) costs; eliminating such bans would reduce (though not eliminate) the scope for collusion. Our finding on rebate bans is consistent with the views of the [Department of Justice \(2005\)](#) as expressed in their complaint in *U.S.A. v. Kentucky Real Estate Commission*; indeed, our findings even accord with surprisingly candid remarks of the real estate agents themselves.^{8,9}

⁶Noting the reluctance of traditional real-estate agents to work with discount brokers, [Levitt and Syverson \(2008a\)](#) argued that it would be in the interests of traditional brokers as a group to refuse to work with discount brokers. However, unlike our work here, their work does not explain why a traditional broker would give up individually profitable opportunities to work with discount brokers (as they assume that each traditional broker internalizes the effect of his decisions on all traditional brokers); moreover, we solve for equilibrium prices while they take the prices of traditional and discount brokers as given.

⁷It is not sufficient in general to consider the repeated version of the reduced normal form game, as the equilibria of that game will not necessarily correspond to subgame-perfect equilibria of the original repeated extensive form game.

⁸The [Department of Justice \(2005\)](#) reports one real estate agent remarking “[A market without rebate bans] would turn into a bidding war, lessen our profits and cheapen our ‘so-called’ profession.” For other remarks by real estate agents, see Appendix A.

⁹The Department of Justice also investigated rebate bans enacted by the Tennessee Real Estate Commission, the South Dakota Real Estate Commission, and the West Virginia Real Estate Commission; all three real estate commissions have rescinded their rebate bans in response to the inquires by the Department of Justice ([Department of Justice and Federal Trade Commission, 2007](#)).

Meanwhile, [Barwick \(2018\)](#) has suggested banning “agency fees,” i.e., the commissions that a seller agent pays to a buyer agent upon the completion of a transaction. As [Barwick \(2018\)](#) noted, countries that have banned agency fees have lower overall commissions (even though the commissions paid by buyers increase). In our framework, the optimal collusive scheme involves fully exploiting sellers while possibly charging buyers less than cost. Eliminating agency fees makes it infeasible to set a buyer fee that is below cost, since agents would no longer be willing to represent buyers. Thus, prices would adjust upwards for buyers but would fall for sellers, and we show that consequently the net agent surplus extracted in the highest-profit equilibrium decreases.

While we focus on real estate agency as our lead application, our analysis can be used to understand other two-sided markets with independent buyer and seller representation. For instance, in the primary market for municipal bonds, spreads in the United States are notoriously high ([Cestau et al., 2019](#)); in that market, underwriters represent municipalities wishing to issue municipal bonds while financial planners and advisors represent high-income individuals wishing to purchase such bonds.

Similarly, both buyers and sellers of stocks listed on the Nasdaq generally worked through dealers. These dealers “protected [their] spreads” (i.e., profits) from so-called “SOES [small order execution system] bandits” by avoiding quoting prices that used odd-eighths of a dollar ([Durr and Colby, 2010](#)) and avoiding transactions with dealers who did quote odd-eighths. This behavior, which was first documented by [Christie and Schultz \(1994\)](#), is closely analogous to the in-period punishments which are key in our model for supporting collusion.¹⁰

The remainder of this paper is organized as follows: Section 2 lays out our model. Section 3 characterizes the optimal collusive prices. Section 4 considers the implications of our work for policy: Section 4.1 considers the effects of rebate bans on prices, Section 4.2 considers the effects of eliminating agency fees on prices, and Section 4.3 considers the extent to which our equilibrium is robust to buyer and seller self-representation. Section 5 concludes.

2 Model

2.1 Framework

We introduce a model of brokered buyer–seller markets. There is a finite set of *agents* A ; we let $\alpha \equiv \frac{1}{|A|}$ be the *market concentration*. Moreover, there is an infinite sequence of unit intervals of short-lived *buyers* $\{B_t\}_{t \in \mathbb{N}}$ and an infinite sequence of unit intervals of short-lived

¹⁰Subsequent to the work of [Christie and Schultz \(1994\)](#), the Securities and Exchange Commission and the Department of Justice investigated the behavior of Nasdaq dealers; related civil litigation led to settlements exceeding \$1 billion dollars ([Christie and Thompson, 2006](#)).

sellers $\{S_t\}_{t \in \mathbb{N}}$. Each agent has a *buyer capacity* κ_B and a *seller capacity* κ_S ; we assume that $\kappa_B \geq \kappa_S$. We require that both κ_B and κ_S are less than $\frac{1}{2}$; that is, no agent can represent more than half of the buyers in any given period, and no agent can represent more than half of the sellers in any given period. We also require that $(|A| - 1)\kappa_B \geq 1$ and $(|A| - 1)\kappa_S \geq 1$; that is, all of the buyers and all of the sellers can be assigned an agent even if one agent is excluded from the economy. Time is discrete and infinite; agents have a common discount factor $\delta \in (0, 1)$.

In each period t , the agents, buyers, and sellers play the following extensive-form stage game:

Step 1: Each agent $a \in A$ offers a *buyer price* $p_{B,t}^a \in \mathbb{R}$ and a *seller price* $p_{S,t}^a \in \mathbb{R}$. All prices are publicly observed.

Step 2: Each buyer $b \in B_t$ submits an ordered list of a subset of agents; that is, each buyer reports a *ranking* over the set of agents, with unlisted agents being unacceptable. Buyers are then assigned via *random rationing* such that no agent is assigned a mass of more than κ_B buyers; this mechanism simply ensures that a buyer prefers to list agents in order of his actual preference and that a buyer is never left unmatched when there is an acceptable agent with unused capacity.¹¹ We denote the agent *representing* buyer b as $\mathbf{a}(b)$, where we let $\mathbf{a}(b) \equiv \emptyset$ if b is unassigned to any agent (i.e., \emptyset represents the *outside option*); moreover, we denote the set of buyers assigned to agent a in period t as $\mathbf{B}_t(a) \equiv \{b \in B_t : \mathbf{a}(b) = a\}$. Similarly (and simultaneously), each seller $s \in S_t$ submits an ordered list of a subset of agents; that is, each seller reports a ranking over the set of agents, with unlisted agents being unacceptable. Sellers are then assigned via *random rationing* such that no agent is assigned a mass of more than κ_S sellers. We denote the agent *representing* seller s as $\mathbf{a}(s)$, where we let $\mathbf{a}(s) = \emptyset$ if s is unassigned to any agent; moreover, we denote the set of sellers assigned to agent a in period t as $\mathbf{S}_t(a) \equiv \{s \in S_t : \mathbf{a}(s) = a\}$. The set of buyers and sellers represented by each agent is publicly observed.

Step 3: Each agent a *invites* a set of other agents. Each invitation includes a contingent (*agency*) fee $f_t^{\tilde{a} \leftarrow a} \in \mathbb{R}$ that will be paid per transaction to the buyer's agent \tilde{a} from the seller's agent a . Each invitation (along with its associated fee) is publicly observed.¹²

¹¹This procedure is the generalization of a random serial dictatorship to settings with a continuum of agents. We formally define the allocation of buyers and sellers to agents given the buyers' and sellers' rankings in Appendix B.

¹²The assumption that an agent observes invitations among other agents is purely for technical convenience; if agents only observe their own invitations, we would have to specify beliefs as a function of invitations received. In particular, in the equilibria we describe, no agent conditions his actions on invitations between agents other than himself.

Step 4: Each agent a *accepts* or *rejects* each invitation that he receives. We denote the set of agents that a invites and that accept a 's invitation, along with a , as $A_t^{\Rightarrow a}$. We denote the set of agents that a accepts invitations from, along with a , as $A_t^{a\Rightarrow}$. After invitations are accepted or rejected, all agents observe which invitations are accepted.

We can think of the set of accepted invitations as generating a (*directed*) *network*, where an edge from \tilde{a} to a represents the fact that \tilde{a} has accepted a 's invitation (and so \tilde{a} 's buyers have access to a 's sellers). We say that the *full network forms among agents* \bar{A} if, for every distinct $\tilde{a}, a \in \bar{A}$, we have that \tilde{a} has accepted a 's invitation (where we use the convention that $\bar{A} = A$ when \bar{A} is unspecified).

Once the network has formed, we model the housing market as a market where each buyer has a unique acceptable seller and each seller has a unique acceptable buyer. Let \mathcal{M}_t be a measurable bijective correspondence between B_t and S_t unknown to all market participants. The correspondence \mathcal{M}_t is a reduced-form way of modeling the preferences of buyers and sellers: for a given buyer b , $\mathcal{M}_t(b) \in S_t$ represents the unique seller whose house is desirable for buyer b . A buyer b will consummate a transaction with a seller s if and only if s 's agent, $\mathbf{a}(s)$, invited b 's agent, $\mathbf{a}(b)$, and $\mathbf{a}(b)$ accepted $\mathbf{a}(s)$'s invitation, i.e., $\mathbf{a}(s) \in A^{\mathbf{a}(b)\Rightarrow}$. We call the fixed value that a buyer receives from a transaction the *buyer surplus* v_B and, similarly, the fixed value that a seller receives from a transaction the *seller surplus* v_S .

Thus, the stage-game payoffs are as follows:

1. The expected payoff for buyer b in period t is given by $(v_B - p_{B,t}^{\mathbf{a}(b)}) |\cup_{a \in A^{\mathbf{a}(b)\Rightarrow}} \mathbf{S}_t(a)|$. That is, the payoff for b is the net value of a purchase for b times the measure of sellers connected to b (which gives the probability of a transaction for b).
2. The expected payoff for seller s in period t is given by $(v_S - p_{S,t}^{\mathbf{a}(s)}) |\cup_{a \in A^{\Rightarrow \mathbf{a}(s)}} \mathbf{B}_t(a)|$. That is, the payoff for s is the net value of a sale for s times the measure of buyers connected to s (which gives the probability of a transaction for s).
3. The expected payoff for an agent a in period t is given by the total revenue from buy-

and sell-side transactions:

$$\begin{aligned}
& \int_{\mathbf{B}_t(a)} \int_S \underbrace{\left(p_{B,t}^a + f_t^{a \leftarrow \mathbf{a}(s)} \right)}_{\text{Transaction profit}} \underbrace{\mathbb{1}_{\{a \in A_t^{\mathbf{a}(s)}\}}}_{\substack{a \text{ invited by } \mathbf{a}(s) \\ \text{and } a \text{ accepted}}} \underbrace{\mathbb{1}_{\{\mathcal{M}_t(b)=s\}}}_{b \text{ and } s \text{ correspond}} ds db \\
& \underbrace{\hspace{10em}}_{\text{Profits as a buy-side agent}} \\
& + \int_{\mathbf{S}_t(a)} \int_B \underbrace{\left(p_{S,t}^a - f_t^{\mathbf{a}(b) \leftarrow a} \right)}_{\text{Transaction profit}} \underbrace{\mathbb{1}_{\{a \in A_t^{\mathbf{a}(b)}\}}}_{\substack{\mathbf{a}(b) \text{ accepted} \\ a \text{'s invitation}}} \underbrace{\mathbb{1}_{\{\mathcal{M}_t(s)=b\}}}_{b \text{ and } s \text{ correspond}} db ds, \\
& \underbrace{\hspace{10em}}_{\text{Profits as a sell-side agent}}
\end{aligned}$$

which reduces to

$$|\mathbf{B}_t(a)| \sum_{\tilde{a} \in A^{\mathbf{a} \Rightarrow}} \left(|\mathbf{S}_t(\tilde{a})| (p_{B,t}^a + f_t^{a \leftarrow \tilde{a}}) \right) + |\mathbf{S}_t(a)| \sum_{\tilde{a} \in A^{\mathbf{a} \Rightarrow}} \left(|\mathbf{B}_t(\tilde{a})| (p_{S,t}^a - f_t^{\tilde{a} \leftarrow a}) \right).$$

2.2 Relationship of Our Assumptions to the Real Estate Industry

While our model is of course stylized, our assumptions are structured to reflect important characteristics for understanding pricing in the real estate market.

Our assumption that there is a unique desirable house for each buyer is designed to capture the idea that preferences among buyers are highly heterogenous. While our assumption is stark, it captures the importance to a buyer of having a large number of houses to consider and to a seller of having a large number of potential buyers.¹³ Adding more realistic preferences for buyers' over houses would significantly complicate the model without adding additional insights; the key assumption is that the expected value to a buyer from hiring a given agent is increasing in the number of sellers linked directly or indirectly to that agent (and, similarly, the expected value to a seller from hiring a given agent is increasing in the number of buyers linked directly or indirectly to that agent).

And while we abstract from the bargaining process between a buyer and seller who have been matched by agents, we could instead allow each matched buyer–seller pair to bargain over the surplus generated by a transaction à la [Hall and Milgrom \(2008\)](#). Then, v_B and v_S represent the (expected) gains for the buyer and seller, respectively, from that bargaining process. Under the assumption that there is a unique desirable house for each buyer, extending the model in this way would have no effect on our results.

Finally, we impose the ordering $\kappa_S \leq \kappa_B$ from our understanding that listing a house for

¹³By contrast, in a commodity market, we would not expect the type of two-sided intermediation studied here to sustain collusion.

sale and the other duties of a selling agent are appreciably more complex and time-consuming than helping a buyer find a suitable house. This ordering is not important for maintaining collusion but does play a role in determining which side of the market pays higher agency fees, as we discuss after Theorem 1.

2.3 Equilibrium Definition

In our setting, perfect collusion could be sustained even in the stage game via coordinated behavior by buyers and sellers; to see this, note that we can support any non-negative prices in a subgame perfect Nash equilibrium of the stage game by having buyers and sellers “coordinate” on not signing up with any agent if some agent deviates on prices. Perfect collusion is thus achieved by having both buyers and sellers refuse to sign up with any agent if prices are not as expected; such behavior is an equilibrium because it is (weakly) optimal for no seller to sign up with any agent if no buyers sign up with any agent and, similarly, it is (weakly) optimal for no buyer to sign up with any agent if no sellers sign up with any agent.¹⁴

The coordination failure equilibria just described are unrealistic in our setting, as they involve a very large number of buyers and sellers coordinating amongst themselves to facilitate collusion by agents. Moreover, a version of the model with a finite number of buyers and sellers who sign up for agents sequentially would not admit such coordination failure equilibria; further, such equilibria would not be robust to allowing firms to offer *insulating tariffs* à la [Weyl \(2010\)](#) and [White and Weyl \(2016\)](#). That is, if brokers can promise to compensate end users if the anticipated number of users on the other side do not show up, coordination failure by end users cannot be used to support collusive equilibria. To avoid such pathological outcomes, we restrict attention to *(buyer-and-seller) coordination-proof equilibria*, which require that no positive mass of buyers and/or sellers can (by altering their actions simultaneously) strictly improve the expected welfare of all of them. Formally, a subgame-perfect Nash equilibrium is *(buyer-and-seller) coordination-proof* if, fixing the strategy profile of the agents, for every period t and every history of play up to t , there does not exist a positive measure subset \bar{B} of buyers and/or positive measure subset \bar{S} of sellers that can, in the agent selection phase, jointly submit different ordered lists that result in higher expected utility for each market participant in $\bar{B} \cup \bar{S}$. Our equilibrium restriction prevents mis-coordination amongst buyers and sellers as a mechanism to support higher prices. Our refinement here is in the spirit of coalition-proofness à la [Bernheim et al. \(1987\)](#) among the short-lived buyers and sellers.¹⁵ Alternatively, one could use a suitably adapted

¹⁴This is a version of the classic result on the indeterminacy of equilibrium in platforms ([Weyl, 2010](#)).

¹⁵Note that the coalition-proofness requirement we are imposing only restricts behavior by buyers and sellers, who are short-lived; thus, we avoid the issues inherent in coalition-proof solution concepts for long-lived

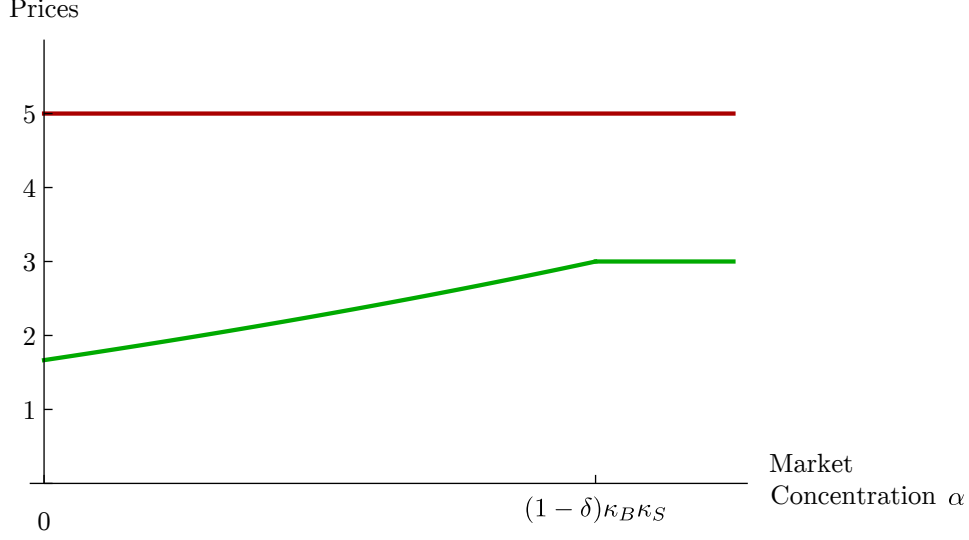


Figure 1: The prices supporting the highest sustainable profit equilibrium. The dark red line is the price agents charge to sellers, p_S^* , and the light green line is the price agents charge to buyers, p_B^* . Here, $\delta = \frac{3}{4}$, $v_B = 3$, $v_S = 5$, $\kappa_B = \frac{1}{5}$, and $\kappa_S = \frac{1}{6}$.

version of the coalitional rationalizability concept introduced by [Ambrus \(2006\)](#) to avoid this type of mis-coordination among short-lived buyers and sellers.

3 Optimal Collusion

We now characterize the highest profits that can collectively be achieved by the agents. We say that a level of total industry profits is *sustainable* if there exists a coordination-proof subgame-perfect Nash equilibrium in which, along the equilibrium path, the total profits obtained by all agents reach that level.

Theorem 1. *For $\delta \geq \frac{1}{2}$, the highest sustainable industry profits are achieved with prices*

$$p_B^* = \begin{cases} v_B & \alpha \geq (1 - \delta)\kappa_B\kappa_S \\ \frac{(1-\delta)\kappa_B(v_B - \kappa_S(v_B + v_S)) + \alpha v_S}{(1-\delta)\kappa_B - \alpha} & \alpha \leq (1 - \delta)\kappa_B\kappa_S \end{cases}$$

$$p_S^* = v_S.$$

Moreover, $\lim_{\alpha \rightarrow 0} (p_B^* + p_S^*) = (v_B + v_S)(1 - \kappa_S) > 0$.

Figure 1 plots the prices offered to buyers and sellers in the highest sustainable profit equilibrium as a function of the market concentration. In this figure, even as market agents in repeated games, as the only agents in our model engaging in coalitional behavior are short-lived. For a discussion of dynamic coalition formation in repeated games, see [Ali and Liu \(2019\)](#).

concentration goes to 0, the seller price remains at the seller surplus v_S ; meanwhile, the buyer price falls (nonlinearly) toward $v_B - \kappa_S(v_B + v_S)$. In general, when the market concentration goes to 0, the buyer price can be positive, i.e., above marginal cost, or negative, i.e., below marginal cost. In Figure 1 the buyer price remains positive for all market concentrations, but it would become negative if the buyer surplus v_B were low enough.

Buyers are offered better prices than sellers in the highest sustainable profits equilibrium, and those prices may in fact be negative. However, the reason for subsidizing buyers is not related to the idea from platform economics that it can be optimal to subsidize one side of the market to encourage adoption by the other side of the market. Rather, buyers are subsidized as such subsidies are the lowest-cost way to discourage agents from undercutting on prices; for this result, it is crucial that buyer capacity κ_B is greater than seller capacity κ_S . If buyer capacity were less than seller capacity, then the equilibrium with the highest sustainable profits would require setting the buyer price to v_B and the seller price below v_S . We now move on to constructing the equilibrium that supports the prices given in Theorem 1, starting with the analysis of the stage game.

3.1 Bertrand Reversion Nash Equilibrium

We first describe the *Bertrand reversion Nash equilibrium* of the stage game. In this equilibrium, each agent announces a buyer price $p_B^a = 0$ and a seller price $p_S^a = 0$. Buyers and sellers then sort themselves equally across agents. Finally, each agent invites each other agent with an agency fee of 0; each agent then accepts every invitation and so the full network forms.

Proposition 1. *There exists a coordination-proof subgame-perfect Nash equilibrium of the stage game in which each agent obtains a payoff of 0, the lowest individually rational payoff. Furthermore, in every symmetric coordination-proof subgame-perfect Nash equilibrium every agent obtains a payoff of 0.*

Working backwards through the stage game, in any subgame-perfect Nash equilibrium each agent a invites every other agent \tilde{a} offering a fee of $-p_B^{\tilde{a}}$, i.e, the fee that captures all of the surplus from any transaction. An agent \tilde{a} accepts any invitation that promises him a non-negative profit, i.e., any invitation with a fee no less than $-p_B^{\tilde{a}}$. We call these actions with respect to making and accepting invitations *statically optimal network formation actions*.

Given statically optimal network formation, if any agent were to be making positive profits, some other agent could undercut him and still make positive profits as buyers and

sellers will now choose to work with the agent with lower prices;¹⁶ thus, competition has the same effect as in standard models of Bertrand competition between firms and drives profits to 0.

3.2 Maintaining Collusion via Network Exclusion

We first provide an intuitive description of a coordination-proof Nash equilibrium that maximizes the surplus extracted by the agents. We then formally construct a strategy profile that delivers prices of (p_B^*, p_S^*) each period, and show that this strategy profile is a coordination-proof subgame-perfect Nash equilibrium. Finally, we show that no other coordination-proof subgame-perfect Nash equilibrium can sustain prices that deliver higher per-period profits than those delivered by (p_B^*, p_S^*) .

The key idea is to construct strategies that incentivize agents not to work with any agent who undercuts the collusive prices. In our equilibrium, play begins in a *cooperation phase*, in which each agent offers a buyer price of p_B^* and a seller price of p_S^* . Assuming all the agents offer p_B^* and p_S^* , buyers and sellers allocate themselves evenly across all agents; agents then form the full network. However, during the cooperation phase, if some agent undercuts on pricing, i.e., becomes a *price deviator*, other agents refuse to form links with him; in light of this, buyers and sellers are less willing to sign on with a price deviator, as they are aware that they will not have as many transaction opportunities when working with such an agent. Hence, if an agent undercuts on price on each side by a small amount, instead of increasing his market share—as one might expect—he finds that no buyer or seller will work with him. Thus, if an agent wants to increase his market share, he must significantly reduce his prices to compensate buyers and sellers for the reduction in transaction opportunities they face for signing up with him; we call the price he must offer to entice buyers the *buyer deviation price (during the cooperation phase)* p_B° and, similarly, the price he must offer to entice sellers the *seller deviation price (during the cooperation phase)* p_S° . We say that (p_B, p_S) is an *effective price deviation* if $(p_B, p_S) \leq (p_B^\circ, p_S^\circ)$.

Of course, to incentivize agents to exclude a price deviator from the network, those agents must expect future rewards from doing so. That is, the “reward should fit the temptation” (Mailath et al., 2017)—and so continuation play must proceed differently depending on whether agents worked with the price deviator, whom we shall refer to as \hat{a} . If all other agents exclude the price deviator \hat{a} from the network, play proceeds to a *collusive punishment phase*. In this phase, prices fall but not to 0; each agent offers a buyer price of q_B^* and a seller price of q_S^* . Moreover, agents continue to exclude the price deviator \hat{a} from the network. During

¹⁶Such intuitive behavior by the buyers and sellers is guaranteed by our assumption of coordination-proofness.

a collusive punishment phase, as in the cooperation phase, a (possibly new) price deviator must substantially undercut (q_B^*, q_S^*) in order to incentivize buyers and sellers to sign up with him; we call the price he must offer to entice buyers the *buyer deviation price (during the collusive punishment phase)* q_B° and, similarly, the price he must offer to entice sellers the *seller deviation price (during the collusive punishment phase)* q_S° . We say that (p_B, p_S) is an *effective price deviation (during the collusive punishment phase)* if $(p_B, p_S) \leq (q_B^\circ, q_S^\circ)$. The prices q_B^* and q_S^* are exactly chosen so that any effective price deviation during the collusive punishment phase is unprofitable.

By contrast, if any agent works with a price deviator (in either the cooperation phase or a collusive punishment phase), play proceeds to a *Bertrand reversion phase*, in which each agent obtains 0 profits in all future periods. Thus, since working with the price deviator leads to 0 profits in all future periods, while not working with the price deviator leads to positive profits in all future periods, sufficiently patient agents will follow through on the threat to exclude a deviator from the network. Note that the degree of patience necessary to incentivize agents to not work with a price deviator does not depend on market concentration: Both future profits from excluding the price deviator and the profits today from working with the price deviator are proportional to $\frac{1}{|A|-1}$, the pro-rated share of each side of the market for each agent other than the price deviator.¹⁷

We now formally construct a strategy profile that sustains (p_B^*, p_S^*) .¹⁸ To simplify the exposition, we first define distinguished actions for buyers and sellers:

1. A buyer (seller) *lists agents arbitrarily* by reporting with equal probability each ranking that includes every agent.
2. A buyer (seller) *prioritizes agent a* by reporting with equal probability each ranking that both includes every agent and ranks a first.
3. A buyer (seller) *deprioritizes agent a* by reporting with equal probability each ranking that both includes every agent and ranks a last.

We further define distinguished actions for agents in the network formation steps:

¹⁷It is unnecessary to revert to Bertrand competition in the event that some agent works with a price deviator. Instead, we could use any equilibrium continuation play that delivers 0 profits to an agent who works with a price deviator; in particular, we could use the collusive punishment phase that punishes an agent who works with a price deviator.

¹⁸Here, we require that buyers and sellers treat identically agents who are treated identically by other agents (with respect to network formation); that is, they do not discriminate between agents who have offered the same prices in this period and whom they expect to form the same network. This restriction prevents implausible coordination by buyers and sellers. In Appendix D, we relax this assumption and show that Theorem 1 still holds. (There, we make a technical simplifying assumption that $\kappa_S \leq \kappa_B \leq \frac{1}{3}$.)

1. The *full network forms with standard fees* when each agent a invites every other agent \tilde{a} with a fee that demands \tilde{a} 's buyer price, i.e., a fee of $-p_B^{\tilde{a}}$, and every agent \tilde{a} accepts every invitation with a fee greater than or equal to $-p_B^{\tilde{a}}$.¹⁹
2. The *network excluding a forms with standard fees* when the full network forms among agents other than a and no agent forms any links with a , that is, when:
 - No agent invites a .
 - Each agent other than a invites every other agent \tilde{a} with a fee that demands \tilde{a} 's buyer price, i.e., a fee of $-p_B^{\tilde{a}}$.
 - Agent a invites every other agent with a fee equal to a 's seller price, p_S^a .
 - Each agent \tilde{a} other than a accepts every invitation he receives with a fee greater than or equal to $-p_B^{\tilde{a}}$ except an invitation from a .
 - Each agent other than a accepts an invitation from a if and only if the fee is (strictly) greater than p_S^a .
 - Agent a accepts an invitation from any other agent if and only if the fee is no less than $-p_B^a$.

We also define the buyer deviation price in the cooperation phase as $p_B^{\circ} = v_B - \frac{1}{\kappa_S}(v_B - p_B^{\star})$ and the seller deviation price $p_S^{\circ} = v_S - \frac{1}{\kappa_B}(v_S - p_S^{\star}) = v_S$; these are the prices at which buyers and sellers will be willing to work with a deviating agent in the cooperation phase. During the collusive punishment phase, prices are $q_B^{\star} = (1 - \kappa_S)v_B - \kappa_S v_S$ and $q_S^{\star} = v_S$. We define the buyer deviation price during the collusive punishment phase as $q_B^{\circ} = v_B - \frac{1}{\kappa_S}(v_B - q_B^{\star})$ and the seller deviation price as $q_S^{\circ} = v_S - \frac{1}{\kappa_B}(v_S - q_S^{\star}) = v_S$; these are the prices at which buyers and sellers will be willing to work with an agent who deviates in the collusive punishment phase. Finally, we do not detail the strategies after mutual deviations, i.e., after two or more agents *simultaneously* deviate: since no agent expects any other agent to deviate, such cases have no effect on incentives.

The strategy profile that sustains p_B^{\star} and p_S^{\star} consists of three phases: In the *cooperation phase*:

1. Each agent offers a buyer price p_B^{\star} and a seller price p_S^{\star} .
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

¹⁹These fees imply that, for a given transaction, the agent representing the seller receives all of the profits obtained by the agents. In fact, it is straightforward to modify the strategies presented here to split the profits more evenly, although doing so requires verifying additional incentive constraints.

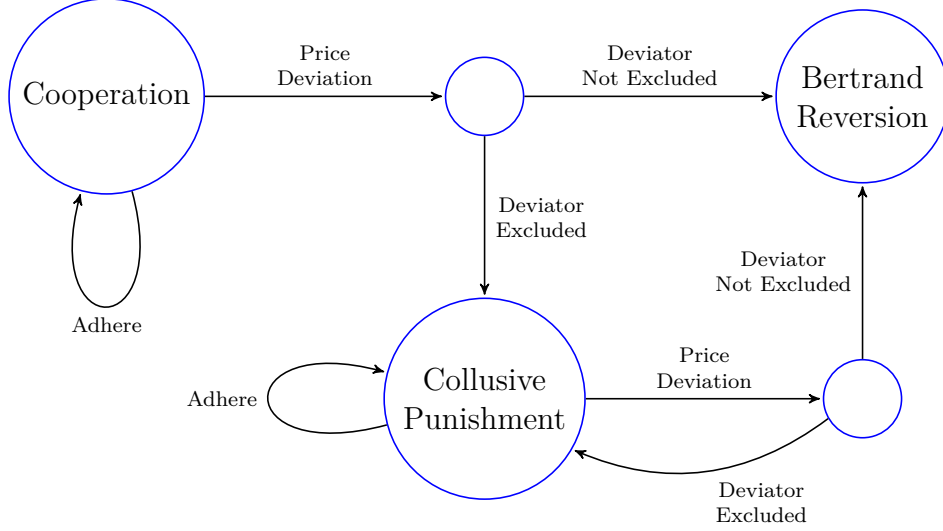


Figure 2: Automaton representation of the equilibrium we consider. Labeled nodes are phases; unlabeled nodes are intermediate phases, which represent the branching of transitions based on behavior in the later steps of the game.

Case 1: *Collusive pricing*: Each agent has offered (p_B^*, p_S^*) . Buyers and sellers list agents arbitrarily.

Case 2: *Ineffective price deviation by \hat{a}* : Each agent except \hat{a} has offered (p_B^*, p_S^*) and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \not\leq (p_B^o, p_S^o)$. Buyers and sellers deprioritize agent \hat{a} .

Case 3: *Effective price deviation by \hat{a}* : Each agent except \hat{a} has offered (p_B^*, p_S^*) and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (p_B^o, p_S^o)$. Buyers and sellers prioritize agent \hat{a} .

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: *Collusive pricing*. The full network forms with standard fees.

Cases 2 and 3: *Price deviation by \hat{a}* . The network excluding \hat{a} forms with standard fees.

4. Under collusive pricing, if the full network with standard fees forms, play continues in the cooperation phase. After a price deviation by \hat{a} , if the network excluding \hat{a} forms with standard fees, then play proceeds to the \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the \hat{a} -collusive punishment phase:

1. Every agent (including \hat{a}) offers a buyer price q_B^* and a seller price q_S^* .
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

Case 1: *Collusive pricing*: Each agent $a \in A$ has offered (q_B^*, q_S^*) . Buyer and sellers deprioritize \hat{a} .

Case 2: *Ineffective price deviation by $\hat{a} \in A$* ²⁰ Each agent $a \in A \setminus \{\hat{a}\}$ has offered (q_B^*, q_S^*) , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \not\leq (q_B^o, q_S^o)$. Buyers and sellers deprioritize \hat{a} .

Case 3: *Effective price deviation by $\hat{a} \in A$* Each agent $a \in A \setminus \{\hat{a}\}$ has offered (q_B^*, q_S^*) , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (q_B^o, q_S^o)$. Buyers and sellers prioritize \hat{a} .

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Cases 1: *Collusive pricing*. The network excluding \hat{a} forms with standard fees.

Cases 2 and 3: *Effective and ineffective price deviations by \hat{a}* . The network excluding \hat{a} forms with standard fees.

4. Under collusive pricing, if the network excluding \hat{a} forms with standard fees, play continues in the \hat{a} -collusive punishment phase. After a price deviation by \hat{a} , if the network excluding \hat{a} forms with standard fees, then play proceeds to the \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

Figure 2 provides an automaton representation of the coordination-proof subgame-perfect Nash equilibrium described here.

It is immediate that the strategy profile just described delivers prices of (p_B^*, p_S^*) in each period; we now verify that it constitutes a coordination-proof subgame-perfect Nash equilibrium.

Making and Responding to Invitations in the Cooperation Phase

We first verify that the prescribed strategy profile is incentive-compatible with respect to making and accepting invitations in the cooperation phase. It is straightforward to show that an agent cannot profitably deviate with respect to making and accepting invitations in the collusive pricing and ineffective price deviation cases. It is also straightforward that, after an effective price deviation by \hat{a} , an agent other than \hat{a} cannot profitably deviate with

²⁰Note that agent \hat{a} could be either \hat{a} or some other agent.

respect to making and accepting invitations from agents other than \hat{a} and that \hat{a} 's prescribed actions are optimal. Details are given in Appendix C.2.

We now analyze a key incentive constraint: that an agent a is better off following his prescribed actions than if he

1. invited \hat{a} with a fee of $-p_B^{\hat{a}}$,
2. accepted an invitation from \hat{a} with a fee of $p_S^{\hat{a}}$, and
3. followed his prescribed actions with respect to other agents.²¹

The payoff for working with the deviator \hat{a} is then given by²²

$$\begin{aligned}
 & \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)}_{\text{Mass of sellers represented by } a} \underbrace{\kappa_B}_{\text{Mass of buyers represented by } \hat{a}} \underbrace{(p_B^{\hat{a}} + p_S^*)}_{\text{Profit per transaction}} + \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_B)}_{\text{Mass of buyers represented by } a} \underbrace{\kappa_S}_{\text{Mass of sellers represented by } \hat{a}} \underbrace{(p_B^* + p_S^{\hat{a}})}_{\text{Profit per transaction}} + \\
 & \hspace{15em} \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)}_{\text{Mass of sellers represented by } a} \underbrace{(1-\kappa_B)}_{\text{Mass of buyers represented by agents other than } \hat{a}} \underbrace{(p_B^* + p_S^*)}_{\text{Profit per transaction}}. \quad (1) \\
 & \hspace{15em} \underbrace{\hspace{15em}}_{\text{Profits from working with agents other than } \hat{a} \text{ this period}}
 \end{aligned}$$

Meanwhile, the total payoff for a from following his prescribed actions is

$$\underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)(1-\kappa_B)(p_B^* + p_S^*)}_{\text{Profits from working with agents other than } \hat{a} \text{ this period}} + \underbrace{\frac{\delta}{1-\delta} \frac{\alpha}{1-\alpha}(q_B^* + q_S^*)}_{\text{Payoff in future periods from adhering}}. \quad (2)$$

Thus, since the profits from working with agents other than \hat{a} this period are identical regardless of whether a works with \hat{a} , it is sufficient that

$$\frac{\delta}{1-\delta} \frac{\alpha}{1-\alpha}(q_B^* + q_S^*) \geq \frac{\alpha}{1-\alpha}(1-\kappa_S)\kappa_B(p_B^{\hat{a}} + p_S^*) + \frac{\alpha}{1-\alpha}(1-\kappa_B)\kappa_S(p_B^* + p_S^{\hat{a}})$$

or, equivalently, that

$$\frac{\delta}{1-\delta}(v_B + v_S)(1-\kappa_S) \geq (1-\kappa_S)\kappa_B(p_B^{\hat{a}} + p_S^*) + (1-\kappa_B)\kappa_S(p_B^* + p_S^{\hat{a}}); \quad (3)$$

²¹It is immediate that a is better off working with the non-deviating agents when working with \hat{a} .

²²Note that $\frac{\alpha}{1-\alpha} = \frac{1}{|A|-1}$.

recall that $q_B^* = (1 - \kappa_S)v_B - \kappa_S v_S$ and $q_S^* = v_S$ and so $q_B^* + q_S^* = (v_B + v_S)(1 - \kappa_S)$. Observe the following:

- Since $p_B^{\hat{a}} \leq p_B^{\circ} \leq p_B^* \leq v_B$ and $p_S^* = v_S$, we have that $p_B^{\hat{a}} + p_S^* \leq v_B + v_S$.
- Since $p_S^{\hat{a}} \leq p_S^{\circ} = p_S^* = v_S$ and $p_B^* \leq v_B$, we have that $p_B^* + p_S^{\hat{a}} \leq v_B + v_S$.
- Since $\kappa_B \geq \kappa_S$, we have that $1 - \kappa_S \geq 1 - \kappa_B$.

Hence, for (3) to hold, it is sufficient that

$$\frac{\delta}{1 - \delta} \geq \kappa_B + \kappa_S;$$

this inequality holds so long as $\delta \geq \frac{1}{2}$ as $\kappa_S \leq \kappa_B \leq \frac{1}{2}$.

Intuitively, each agent a other than \hat{a} is unwilling to work with the price deviator since the number of buyers and sellers represented by a this period is roughly proportional to α and future profits for a are also roughly proportional to α . Thus, for reasonably high discount factors, future profits for a are worth more than the gains from working with \hat{a} today.

Making and Responding to Invitations in the Collusive Punishment Phase

The analysis of making and responding to invitations in the \hat{a} -collusive punishment phase proceeds very similarly to the analysis of making and responding to invitations in the cooperation phase. Details are given in Appendix C.2.

Agent Selection in the Cooperation Phase

There are three cases to consider:

Case 1: Collusive Pricing. It is straightforward that no positive masses of buyers and sellers can alter their actions to simultaneously improve their welfare since every agent offers a buyer price $p_B^* \leq v_B$ to buyers, every agent offers a seller price $p_S^* \leq v_S$ to sellers, every buyer and seller obtains an agent, and the full network forms (regardless of buyer and seller actions).

Cases 2 and 3: Effective and ineffective price deviations by \hat{a} . If \hat{a} offers prices other than (p_B^*, p_S^*) , both buyers and sellers anticipate that \hat{a} will not have invitations accepted or accept any invitations. Thus, any buyer who is represented by \hat{a} will only see sellers represented by \hat{a} and, similarly, any seller who is represented by \hat{a} will only see buyers

represented by \hat{a} . Hence, a positive mass of buyers μ_B and a positive mass of sellers μ_S will both be strictly better off when represented by \hat{a} if and only if both

$$\underbrace{(v_B - p_B^{\hat{a}})}_{\text{Buyer } b\text{'s payoff from a transaction}} \underbrace{\mu_S}_{\text{Probability of a transaction for } b} \geq v_B - p_B^* \quad \text{and} \quad \underbrace{(v_S - p_S^{\hat{a}})}_{\text{Seller } s\text{'s payoff from a transaction}} \underbrace{\mu_B}_{\text{Probability of a transaction for } s} \geq v_S - p_S^*;$$

that is, if both

$$p_B^{\hat{a}} \leq v_B - \frac{1}{\mu_S}(v_B - p_B^*) \quad \text{and} \quad p_S^{\hat{a}} \leq v_S - \frac{1}{\mu_B}(v_S - p_S^*).$$

These conditions are easiest to satisfy when $\mu_B = \kappa_B$ and $\mu_S = \kappa_S$; thus, buyers and sellers will work with \hat{a} so long as both

$$p_B^{\hat{a}} \leq v_B - \frac{1}{\kappa_S}(v_B - p_B^*) = p_B^{\circ} \quad \text{and} \quad p_S^{\hat{a}} \leq v_S - \frac{1}{\kappa_B}(v_S - p_S^*) = p_S^{\circ}.$$

Intuitively, buyers demand a large discount to work with a since a , by lowering his buyer price, has effectively become a *lower quality* agent because other agents no longer work with a .

Agent Selection in the Collusive Punishment Phase

The analysis of agent selection in the \hat{a} -collusive punishment phase is analogous to that for agent selection in the cooperation phase. Details are given in Appendix C.2.

Deviating on Prices in the Collusive Punishment Phase

We now check that no agent has an incentive to post prices at or below $(q_B^{\circ}, q_S^{\circ})$ in the \hat{a} -collusive punishment phase. If an agent does post prices at or below $(q_B^{\circ}, q_S^{\circ})$, he obtains his full capacity of buyers and sellers but does not work with any other agents. Thus, his profits are given by

$$\begin{aligned} \kappa_B \kappa_S (q_B^{\circ} + q_S^{\circ}) &= \kappa_B \kappa_S \left(\left(v_B - \frac{1}{\kappa_S}(v_B - q_B^*) \right) + v_S \right) \\ &= \kappa_B \kappa_S \left(v_B - \frac{1}{\kappa_S}(v_B - (1 - \kappa_S)(v_B - \kappa_S v_S)) + v_S \right) \\ &= \kappa_B \kappa_S \left(v_B - \frac{1}{\kappa_S}(\kappa_S v_B + \kappa_S v_S) + v_S \right) \\ &= 0. \end{aligned}$$

Since no agent (including \hat{a}) receives less than 0 in the collusive punishment phase, no agent will deviate.

Deviating on Prices in the Cooperation Phase

Finally, we verify that, during the cooperation phase, no agent has an incentive to deviate on prices. Along the equilibrium path, an agent's profits are given by $\frac{1}{1-\delta}\alpha(p_B^* + p_S^*)$. Following the maximal effective price deviation (p_B°, p_S°) , an agent's profits are given by $(p_B^\circ + p_S^\circ)\kappa_B\kappa_S$. Thus, prescribed play is optimal if

$$\frac{\alpha}{1-\delta}(p_B^* + p_S^*) \geq (p_B^\circ + p_S^\circ)\kappa_B\kappa_S;$$

that is, so long as

$$\frac{\alpha}{1-\delta}(p_B^* + p_S^*) \geq \left(\left(v_B - \frac{1}{\kappa_S}(v_B - p_B^*) \right) + \left(v_S - \frac{1}{\kappa_B}(v_S - p_S^*) \right) \right) \kappa_B\kappa_S. \quad (4)$$

Hence, the highest sustainable profits are found by solving

$$\max_{p_B^*, p_S^*} \{p_B^* + p_S^*\}$$

subject to (4) and the individual rationality constraints for the buyers and sellers ($p_B^* \leq v_B$ and $p_S^* \leq v_S$). Solving this linear program yields

$$p_B^* = \min \left\{ \frac{(1-\delta)\kappa_B(v_B - \kappa_S(v_B + v_S)) + \alpha v_S}{(1-\delta)\kappa_B - \alpha}, v_B \right\}$$

$$p_S^* = v_S.$$

Maximality of Surplus Extraction

Finally, we need to show that no surplus extraction greater than $p_B^* + p_S^* = p_B^* + v_S$ can be maintained. If p_B^* is exactly v_B , then the individual rationality of buyers and sellers ensures that no greater surplus extraction is possible. If p_B^* is less than v_B , then any prices (p'_B, p'_S) that generate more surplus extraction than (p_B^*, p_S^*) must distribute that surplus so that at least one agent a receives no more than $\alpha(p'_B + p'_S)$. But a could then increase his total profits by choosing sufficiently low prices to attract buyers and sellers who understand that they will only have access to sellers and buyers respectively who are also represented by a .

3.3 Discussion

3.3.1 Supporting Collusion with Grim Trigger Strategies

Our results show prices appreciably above marginal cost can be sustained regardless of the number of agents. To compare our results with the standard analysis of Bertrand competition games, we consider here whether simpler “grim trigger” strategies will also allow agents to extract full surplus. The grim trigger strategy profile that obtains full surplus extraction involves every agent offering a buyer price of v_B and a seller price of v_S , statically optimal network formation (regardless of offered prices), and any deviation leading to Bertrand competition in future periods. If agents are playing a grim trigger strategy profile, then an optimal deviation involves an ϵ price cut to buyers and sellers and statically optimal network formation actions; this generates a profit of $(v_B + v_S)\kappa_S$ for the deviating agent. Adhering to a grim trigger strategy generates a profit of $\frac{\alpha}{1-\delta}(v_B + v_S)$. Thus, grim trigger strategies will only be effective if the market concentration α is higher than $(1 - \delta)\kappa_S$; hence, market concentration must be quite high to maintain collusion using grim trigger strategies. Moreover, such strategies cannot facilitate any collusion when market concentration is below $(1 - \delta)\kappa_S$. By contrast, the strategy profile described in Section 3.2 allows full surplus extraction by the agents so long as α is greater than $(1 - \delta)\kappa_B\kappa_S$ (which is less than $(1 - \delta)\kappa_S$); that strategy profile continues to facilitate significant surplus extraction even as the number of agents grows large (i.e., when we might naïvely expect the market to become “perfectly competitive”).

3.3.2 Characteristics of Brokered Markets Crucial for Collusion

There are two characteristics of the model that are crucial for the existence of highly collusive behavior even when the number of players is large. First, both buyers and sellers obtain independent representation but expect their representative to work with the representatives of others; this structure allows for in-period punishments of the type we describe.²³ Adding a restriction that buyers and sellers can only meet when they use the same agent switches the model to a traditional platform model with limited capacity; under this alternative structure, collusion can no longer be maintained at low discount factors. Such a change, however, reduces total welfare and can even lower consumer surplus since preventing network formation leads to an inefficient lack of transactions in the market; each buyer has at most a κ_S probability of buying a house, and each seller has at most a κ_B probability of selling a house. If κ_S and κ_B are small, few transactions will be completed, reducing the gains from

²³In fact, in many other countries such as the United Kingdom and Australia, buyers typically do not obtain independent representation, and fees for real estate transactions are considerably lower; see [Delcours and Miller \(2002\)](#).

competitive pricing.

Second, the effect of κ_S and κ_B on the highest sustainable price highlights the other feature of our model that allows for high prices in an unconcentrated industry: Matching buyers and sellers is challenging since there is only one potential buyer for each house and only one acceptable house for each buyer. While extreme, these assumptions capture the importance of match quality in the housing market. If, by contrast, each house were acceptable to a large number of buyers and each buyer found many houses acceptable, collusion would be much more difficult to sustain; in this case, an entrant offering service to both sides of the market need only offer slightly lower prices to attract both buyers and sellers, as a buyer (or seller) would still have a high probability of finding an acceptable match even if the entrant were excluded from the network. Thus, when buyers and sellers have a large number of possible transaction partners, the price for facilitating a transaction should be close to marginal cost.

Examples of intermediated markets where we do not seem to see collusion include ridesharing firms and travel agents selling airline tickets. However, in both of these markets, intermediaries have sufficient capacity to serve all of one side of the market: Ridesharing firms have effectively infinite capacity on both sides of the market, while a typical travel agent effectively represents all airlines. In fact, in both the ridesharing and travel agency markets, sellers “multi-home” and typically work with each platform intermediary; thus, there is no need for platform intermediaries to form a network, and so our analysis does not apply. Even if ridesharing firms exhibited limited capacity on both sides of the market, the fact that match quality does not matter very much would imply that the ability to punish other firms via network exclusion would be very weak; so even in this case, we would not expect firms to be able to sustain high prices.

3.3.3 Sharing Price Information

A key requirement for successful collusion in our setting is that agents are able to observe other agents’ pricing on both the buy- and sell-sides. Consistent with this, a number of features of the industry support mandatory price disclosure: First, realtors are supposed to list any rebate to buyers on the HUD-1 Settlement Statement.²⁴ Second, we frequently see agents aggregating into brokerages that share infrastructure and administrative resources and often feature some amount of revenue sharing as well; while this arrangement may reduce costs, it also promotes the observability and coordination of pricing. In particular, real estate brokerages both create and enforce norms around pricing among their own agents, which

²⁴See, e.g., the advice given by the Texas Association of Realtors at <https://www.texasrealestate.com/members/legal-and-ethics/resources/legal-faq/sharing-fees/>.

facilitates overall price observability.²⁵

However, we do not need prices to be perfectly observable in order to facilitate collusion. If there was just a small probability that other agents will not detect a price deviation, there still exist equilibria in which the prices charged to buyers and sellers are bounded away from marginal costs even as the number of agents becomes arbitrarily large. One form of these equilibria requires designating a fixed number of agents as “insiders” who each obtain a small but non-vanishing share of surplus each period, while “outsiders” are largely excluded from the market. Insiders do not cut prices since the loss of future profits is greater than the expected gains from cutting prices today, and prices are chosen (as in the collusive punishment phase of our main equilibrium) so that no outsider can attract buyers and sellers while still making a positive profit.

3.3.4 Entry by a Two-Sided Platform

The dynamics we have identified here can also help explain why entry, including entry by discount realtors and the like, has been unable to drive prices down to competitive levels; this is particularly surprising in light of the fact that the price of many other financial transactions has collapsed over the past few decades—e.g., stock commissions and mutual fund fees. Our analysis shows that, in equilibrium, such entrants can be “frozen out” by market incumbents who refuse to work with them. Likewise, our analysis can help explain why two-sided real estate platforms (e.g., Redfin) have so far not broken the grip of traditional realtors.²⁶ To get around being excluded, the platform entrant must, at a minimum, be quite large (in the sense of having high capacity to serve both home buyers and home sellers), since otherwise it will not prove valuable to customers because it will lack access to home buyers and home sellers not on the platform. But most entrants cannot achieve high capacity at launch, especially in a two-sided market such as real estate, in which a firm must provide services to both buyers and sellers. Moreover, achieving high capacity in real estate brokerage services is particularly difficult, as expanding capacity means hiring new employees, and those employees must possess real estate licenses as well as considerable implicit local knowledge.

²⁵See, for example, the case of *U.S. v. Foley et al.*, 598 F.2d 1323 (4th Cir. 1979) mentioned in Footnote 2, in which senior members of various real estate brokerages discussed shifting pricing from 6% to 7%, and promised that their agents would go along with the change.

²⁶Note that real estate listing websites such as Zillow and Trulia are *not* competitors to real estate agents; indeed, these websites are common venues for real estate agent advertising.

4 Policy Responses

4.1 Rebate Bans

In the context of real estate regulation, a “rebate ban” prohibits a buyer’s agent from paying a buyer for the right to represent him in a real estate transaction—that is, a rebate ban prohibits a buyer’s agent from “rebating” part of the commission that he receives from the seller’s agent back to the buyer. Such rebate bans are a feature of real estate regulation in several U.S. states (Magura, 2007). Han and Hong (2011) argued that rebate bans are anticompetitive and estimated a model in which rebate bans generate excessive entry into real estate brokerage.

We show that rebate bans can play an important role in enhancing the profitability of collusive behavior by agents. We model a *rebate ban* as a constraint that each agent must offer a weakly positive buyer price, i.e., that $p_{B,t}^a \geq 0$ for all agents $a \in A$ and all times t . Rebate bans can facilitate collusion even when no agent is offering a 0 price to buyers, as rebate bans constrain off-equilibrium path pricing; for a given collusive pricing scheme, a deviator may need to charge a negative price to attract buyers, as this is the only way to sufficiently compensate buyers for their reduced access to sellers. Thus, when negative buyer prices are prohibited, it is harder for the deviator to recruit buyers and so higher prices for buyers (and higher profits) can be sustained.

Theorem 2. *For $\delta \geq \frac{1}{2}$, the highest sustainable industry profits in a coordination-proof subgame-perfect Nash equilibrium with a rebate ban are achieved with prices*

$$p_B^* = \begin{cases} v_B & \hat{\alpha} \geq \kappa_S \\ \frac{(v_B - \kappa_S(v_B + v_S)) + \hat{\alpha}v_S}{1 - \hat{\alpha}} & \hat{\alpha} \in \left[\kappa_S \frac{v_S}{v_S + (1 - \kappa_S)v_B}, \kappa_S \right] \\ v_B(1 - \kappa_S) & \hat{\alpha} \leq \kappa_S \frac{v_S}{v_S + (1 - \kappa_S)v_B} \end{cases}$$

$$p_S^* = v_S,$$

where $\hat{\alpha} = \frac{\alpha}{(1 - \delta)\kappa_B}$; profits are strictly higher than without a rebate ban whenever $\hat{\alpha} < \kappa_S \frac{v_S}{v_S + (1 - \kappa_S)v_B}$.

Moreover, $\lim_{\alpha \rightarrow 0} (p_B^* + p_S^*) = v_B(1 - \kappa_S) + v_S$.

Figure 3 plots the prices offered to buyers and sellers in the highest sustainable profit equilibrium as a function of the market concentration when a rebate ban is present. Intuitively, rebate bans only have an effect when the optimal deviating buyer price p_B^o is negative; this happens when market concentration is sufficiently low, i.e., $\hat{\alpha} < \kappa_S \frac{v_S}{v_S + (1 - \kappa_S)v_B}$. Thus, when market concentration is sufficiently high, i.e., $\hat{\alpha} \geq \kappa_S \frac{v_S}{v_S + (1 - \kappa_S)v_B}$, the rebate ban has no




Figure 3: The prices supporting the highest sustainable profit equilibrium when a rebate ban is present. The dark red line is the seller price, p_S^* , and the light green line is the buyer price, p_B^* . The light green dashed line is the buyer price when no rebate ban is present; the seller price is invariant to the presence of a rebate ban. Here, $\delta = \frac{3}{4}$, $v_B = 3$, $v_S = 5$, $\kappa_B = \frac{1}{5}$, and $\kappa_S = \frac{1}{6}$.

effect. By contrast, when market concentration is sufficiently low, i.e., $\hat{\alpha} < \kappa_S \frac{v_S}{v_S + (1 - \kappa_S)v_B}$, the collusive buyer price p_B^* is set so that buyers would find working with an agent with a buyer price of 0 (and only getting access to that agent's sellers) to be weakly dispreferred to working with an agent with a buyer price of p_B^* (and getting access to all sellers).

4.2 Agent Specialization and Agency Fees

In our baseline model, all agents are *ex ante* identical, and we show that a symmetric equilibrium delivers the highest sustainable surplus extraction. Within the real estate industry, however, there is at least some distinction between agents who primarily serve buyers and agents who primarily serve sellers. In particular, while any agent can and likely does represent buyers at times, there are agents, and even firms of agents, who exclusively represent buyers.²⁷

In the U.S. the price paid by a buyer to his agent is typically zero, despite some costs for the agent of representing a buyer; indeed, this is consistent with our model, which suggests that optimal collusion might require agents to subsidize buyers. Thus, it follows that agent specialization may affect the structure of optimal collusion. Allowing for different types of agents may be important, as our model predicts that optimal collusion may require that

²⁷See, for example, the National Association of Exclusive Buyer Agents (<http://naeba.org/>).

agents subsidize buyers rather than charge buyers for access to the platform.

To capture the potential for different roles among agents, we generalize our model to allow for two types of agents. *Seller-proficient agents* have a low cost for acting as a seller's agent and a potentially high cost for acting as a buyer's agent. *Buyer-exclusive agents* have a low cost for acting as a buyer's agent and a prohibitively high cost for acting as a seller's agent. The asymmetry in our setup (seller proficiency versus buyer exclusivity) is in line with the industry structure: while there are few real estate brokers or agents who represent only sellers, there do exist buyer-exclusive agencies, and even a trade association for them, the National Association of Exclusive Buyer Agents. As we show, agent heterogeneity can play an important role in the structure of optimal collusion. In particular, buyer-exclusive agents cannot be induced to set negative buyer prices unless they are compensated through agency fees. Thus, our extension of the model allows us to investigate the effects of eliminating agency fees, a reform originally suggested by [Barwick \(2018\)](#).

4.2.1 Model with Buyer-Exclusive and Seller-Proficient Agents

We modify our model by partitioning the set of agents A into a set of buyer-exclusive agents A_B and a set of seller-proficient agents A_S ; we let $\beta \equiv \frac{1}{|A_B|}$ and $\sigma \equiv \frac{1}{|A_S|}$. We also introduce the cost parameter c as the cost a seller-proficient agent incurs when he represents a buyer in a transaction; to state our results more simply, we assume that $c \leq v_B$. The cost for buyer-exclusive agents to represent a seller is assumed to be high enough to make such an action clearly undesirable; for simplicity we assume that buyer-exclusive agents cannot represent sellers and so only offer a buyer price. Seller-proficient agents are allowed, but not required, to offer a buyer price. Consistent with the symmetric model, we normalize the cost of buyer-exclusive agents representing buyers and seller-proficient agents representing sellers to 0. Thus, the stage game payoff to a buyer-exclusive agent a is

$$|\mathbf{B}_t(a)| \sum_{\tilde{a} \in A^{a \Rightarrow}} \left(|\mathbf{S}_t(\tilde{a})| (p_{B,t}^a + f_t^{a \leftarrow \tilde{a}}) \right),$$

while the payoff to a seller-proficient agent is

$$|\mathbf{B}_t(a)| \sum_{\tilde{a} \in A^{a \Rightarrow}} \left(|\mathbf{S}_t(\tilde{a})| (p_{B,t}^a + f_t^{a \leftarrow \tilde{a}} - c) \right) + |\mathbf{S}_t(a)| \sum_{\tilde{a} \in A^{\Rightarrow a}} \left(|\mathbf{B}_t(\tilde{a})| (p_{S,t}^a - f_t^{\tilde{a} \leftarrow a}) \right).$$

4.2.2 Optimal Collusion

We now characterize the highest profits that can be collectively achieved by the industry.




Figure 4: The prices supporting the highest sustainable profit equilibrium. The dark red line is the seller price, p_S^* , and the light green line is the buyer price (when agency fees are allowed). Here, $\delta = \frac{3}{4}$, $v_B = 1$, $v_S = 10$, $\kappa_B = \frac{1}{4}$, $\kappa_S = \frac{1}{5}$, and $c = 1$.

Theorem 3. *For $\delta \geq \frac{1}{2}$, the highest sustainable industry profits (and the highest sustainable profits for seller-proficient agents) are achieved with prices*

$$p_B^* = \begin{cases} v_B & \hat{\sigma} \geq \kappa_S \left(1 - \frac{c}{v_B + v_S}\right) \\ \frac{v_B - \kappa_S(v_B + v_S - c) + \hat{\sigma}v_S}{1 - \hat{\sigma}} & \hat{\sigma} \leq \kappa_S \left(1 - \frac{c}{v_B + v_S}\right) \end{cases}$$

$$p_S^* = v_S$$

where $\hat{\sigma} = \frac{\sigma}{(1-\delta)\kappa_B}$.

Moreover, $\lim_{\sigma \rightarrow 0} (p_B^* + p_S^*) = (v_B + v_S)(1 - \kappa_S) + \kappa_S c$, which is strictly higher than the highest sustainable industry profits without agent specialization.

It is immediate from the statement of the theorem that prices are higher under specialization. Prices are higher because a two-sided deviation is less profitable for seller-proficient agents (and impossible for buyer-exclusive agents); meanwhile, no agent incurs the additional cost c along the equilibrium path, as each agent only works with one side of the market. In fact, if working with buyers is sufficiently costly for seller-proficient agents (i.e., $c \geq v_B + v_S$), then monopoly prices can be sustained. Equilibrium prices are exhibited in Figure 4; note that when market concentration is low enough, buyers are charged negative prices.

The overall structure of the equilibrium strategy profile that supports these prices is similar to the strategy profile that supports the prices given in Theorem 1. The key difference in equilibrium behavior is that only buyer-exclusive agents represent buyers and only seller-proficient agents represent sellers. Additionally, we must now separately verify

that buyer-exclusive and seller-proficient agents will not work with a seller-proficient agent following his price deviation. In fact, to provide the most effective dynamic incentives, we need two different collusive punishment phases: The first collusive punishment phase follows a price deviation in which the price deviator attracts both buyers and sellers; this phase uses prices (q_B^*, q_S^*) and an agency fee g^* . The second collusive punishment phase follows a price deviation in which the deviator attracts only sellers; this phase uses the same prices (q_B^*, q_S^*) but a new agency fee h^* . In both cases, the prices (q_B^*, q_S^*) are given by

$$\begin{aligned} q_S^* &= v_S \\ q_B^* &= v_B - \kappa_S(v_B + v_S) + \kappa_S c. \end{aligned}$$

Note that we do not need to worry about price deviations by either a buyer-exclusive or a seller-proficient agent \hat{a} on the buy-side; for such a deviation, each seller-proficient agent a (still) invites \hat{a} with a fee that obtains all of the surplus from a transaction between \hat{a} 's buyers and a 's sellers.

The full proof is relegated to the appendix, but here we describe the main differences in the analysis:

First, we show that, given collusive punishment phases with prices (q_B^*, q_S^*) , following a price deviation by some seller-proficient agent \hat{a} in the cooperation phase, no agent will work with \hat{a} within-period. There are two relevant price deviations to consider: in the first, \hat{a} attracts both buyers and sellers, while in the second, \hat{a} just attracts sellers.

When \hat{a} attracts both buyers and sellers, we enter the Bertrand reversion phase if any agent accepts an invitation from \hat{a} ; if no invitation is accepted we enter a collusive punishment phase with prices $(q_B^*, q_S^*) = (v_B - \kappa_S(v_B + v_S) + \kappa_S c, v_S)$ and an agency fee g^* . Note that \hat{a} will never make an invitation with a fee greater than $p_S^{\hat{a}} \leq p_S^*$; thus, to incentivize buyer-exclusive agents to reject invitations from \hat{a} it is sufficient that

$$\underbrace{\frac{\delta}{1-\delta} \beta (q_B^* + g^*)}_{\text{Continuation payoff for a buyer-proficient agent for excluding } \hat{a}} \geq \underbrace{\beta (1 - \kappa_B) \kappa_S (p_B^* + p_S^*)}_{\text{Gain to a buyer-exclusive agent for working with } \hat{a}}.$$

Moreover, \hat{a} will never accept an invitation with a fee less than $c - p_B^{\hat{a}} \geq c - p_B^*$; thus, to incentivize seller-proficient agents to not make invitations to \hat{a} it is sufficient that

$$\underbrace{\frac{\delta}{1-\delta} \frac{\sigma}{1-\sigma} (q_S^* - g^*)}_{\text{Continuation payoff for a seller-proficient agent for excluding } \hat{a}} \geq \underbrace{\kappa_B \frac{\sigma}{1-\sigma} (1 - \kappa_S) (p_S^* + p_B^* - c)}_{\text{Gain to a seller-proficient agent from working with } \hat{a}}.$$

To show the existence of a g^* that simultaneously satisfies both inequalities, it is sufficient that

$$\frac{\delta}{1-\delta}(q_B^* + q_S^*) \geq (1 - \kappa_B)\kappa_S(p_B^* + p_S^*) + \kappa_B(1 - \kappa_S)(p_S^* + p_B^* - c).$$

Thus, as $q_B^* + q_S^* \geq (v_B + v_S)(1 - \kappa_S)$, $p_B^* + p_S^* \leq v_B + v_S$, $\kappa_S \leq \kappa_B \leq \frac{1}{2}$, and $c \geq 0$, it is sufficient that

$$\frac{\delta}{1-\delta}(v_B + v_S)(1 - \kappa_S) \geq (1 - \kappa_B)\kappa_S(v_B + v_S) + \kappa_B(1 - \kappa_S)(v_B + v_S),$$

which holds if

$$\frac{\delta}{1-\delta} \geq \kappa_S + \kappa_B,$$

and so it is enough that $\delta \geq \frac{1}{2}$. Note that any g^* which satisfies both inequalities will be in $[-q_B^*, q_S^*]$.

When \hat{a} attracts only sellers, we enter the Bertrand reversion phase if any buyer-exclusive agent accepts an invitation from \hat{a} ; if no invitation is accepted we enter a collusive punishment phase with prices $(q_B^*, q_S^*) = (v_B - \kappa_S(v_B + v_S) + \kappa_S c, v_S)$ and an agency fee h^* . Note that \hat{a} will never make an invitation with a fee greater than $p_S^{\hat{a}} \leq p_S^*$; thus, to incentivize buyer-exclusive agents to reject invitations from \hat{a} it is sufficient that

$$\underbrace{\frac{\delta}{1-\delta}\beta(q_B^* + h^*)}_{\text{Continuation payoff for a buyer-proficient agent for excluding } \hat{a}} \geq \underbrace{\beta\kappa_S(p_B^* + p_S^*)}_{\text{Gain to a buyer-exclusive agent from working with } \hat{a}}.$$

Letting $h^* = q_S^*$, and recalling that $q_B^* + q_S^* \geq (v_B + v_S)(1 - \kappa_S)$, $p_B^* + p_S^* \leq v_B + v_S$, and $\kappa_S \leq \frac{1}{2}$, it is sufficient that

$$\frac{\delta}{1-\delta}\beta(v_B + v_S)(1 - \kappa_S) \geq \beta\kappa_S(v_B + v_S).$$

which holds if

$$\frac{\delta}{1-\delta} \geq \frac{\kappa_S}{1 - \kappa_S}$$

and so it is enough that $\delta \geq \frac{1}{2}$. Note that, while seller-proficient agents will receive zero surplus in the ensuing collusive punishment phase, the prices q_B^* and q_S^* were chosen so that any price deviation—even one that attracts both buyers and sellers—will be unprofitable.

Second, the same pair of collusive punishment phases will deter other agents from working with an agent who deviates on price in the collusive punishment phase. This result follows analogously to the preceding analysis showing our collusive punishment phases will deter agents from working with an agent who deviates on price in the cooperation phase.

Third, we need to determine the highest prices (p_B°, p_S°) a price deviator can offer in the cooperation phase that will successfully attract buyers and sellers. As in the analysis of our baseline model, we calculate that the highest prices a price deviator can offer—and still attract buyers and sellers—as a function of (p_B^*, p_S^*) ; these prices are given by

$$p_B^\circ = v_B - \frac{1}{\kappa_S}(v_B - p_B^*) \text{ and } p_S^\circ = v_S - \frac{1}{\kappa_B}(v_S - p_S^*).$$

Fourth, and similarly, the best prices a price deviator can offer—and still attract buyers and sellers—in the collusive punishment phase as a function of (q_B^*, q_S^*) are given by

$$q_B^\circ = v_B - \frac{1}{\kappa_S}(v_B - q_B^*) \text{ and } q_S^\circ = v_S - \frac{1}{\kappa_B}(v_S - q_S^*).$$

Fifth, we show that we can sustain the agency fees g^* and h^* corresponding to the two different collusive punishment phases. To do this, we require that a buyer-exclusive agent reject any agency fee less than g^* ; if any agency fee less than g^* is rejected, then it will be incentive-compatible for seller-proficient agents to offer agency fees of g^* . If a buyer-exclusive agent rejects an agency fee less than g^* , future play continues in the same collusive punishment phase, and so his profits will be

$$\frac{\delta}{1 - \delta} \beta (q_B^* + g^*);$$

while his profits from working with a seller-proficient agent offering a fee less than g^* are at most

$$\beta \frac{\sigma}{1 - \sigma} (q_B^* + g^*).$$

Thus, since $\sigma \leq \frac{1}{2}$ (i.e., there are at least two seller-proficient agents), it is sufficient that $\delta \geq \frac{1}{2}$ and $g^* \geq -q_B^*$. We also require that a seller-proficient agent (other than \hat{a}) will offer agency fees of g^* ; this is incentive-compatible so long as $g^* \leq q_S^*$.

We now show we can sustain $h^* = q_S^* = v_S$ in a collusive punishment phase with prices (q_B^*, q_S^*) . To do this, we require that a buyer-exclusive agent reject any agency fee less than q_S^* ; if any agency fee less than q_S^* is rejected, then it will be incentive-compatible for seller-proficient agents to offer agency fees of q_S^* . If a buyer-exclusive agent rejects an agency fee less than q_S^* , future play continues in the same collusive punishment phase, and so his

profits will be

$$\frac{\delta}{1-\delta}\beta(q_B^* + h^*);$$

while his profits from working with a seller-proficient agent offering a fee less than q_S^* are at most

$$\beta\frac{\sigma}{1-\sigma}(q_B^* + h^*).$$

Thus, since $\sigma \leq \frac{1}{2}$ (i.e., there are at least two seller-proficient agents), it is sufficient that $\delta \geq \frac{1}{2}$.

Sixth, we show that the collusive punishment prices of (q_B^*, q_S^*) can be sustained. Here, it is necessary to check that no seller-proficient agent will wish to offer prices sufficiently low to attract buyers and sellers; in particular, we need to ensure the punished seller-proficient agent, who receives 0 in the collusive punishment phase, will not wish to offer prices sufficiently low to attract buyers and sellers. Thus, we need that

$$\begin{aligned} \kappa_B\kappa_S(q_B^\circ + q_S^\circ - c) &= \kappa_B\kappa_S\left(\left(v_B - \frac{1}{\kappa_S}(v_B - q_B^*)\right) + \left(v_S - \frac{1}{\kappa_B}(v_S - q_S^*)\right) - c\right) \\ &= \kappa_B\kappa_S\left(\left(v_B - \frac{1}{\kappa_S}(v_B - (v_B - \kappa_S(v_B + v_S) + \kappa_S c))\right) + v_S - c\right) \\ &= \kappa_B\kappa_S\left(\left(v_B - \frac{1}{\kappa_S}(\kappa_S v_B + \kappa_S v_S - \kappa_S c) + v_S - c\right)\right) \\ &= 0. \end{aligned}$$

That is, we have chosen (q_B^*, q_S^*) to provide the highest profits possible while ensuring that the punished agent cannot obtain positive profits.

Seventh, and finally, we derive the most profitable pair of prices such that no seller-proficient agent has an incentive to offer prices of (p_B°, p_S°) (thus attracting buyers and sellers) and work alone. Thus, the highest sustainable profits are given by

$$\max_{p_B^*, p_S^*} \{p_B^* + p_S^*\}$$

subject to

$$\begin{aligned} \frac{1}{1-\delta}\sigma(p_B^* + p_S^*) &\geq \kappa_B\kappa_S(p_B^\circ + p_S^\circ - c) \\ p_B^* &\leq v_B \\ p_S^* &\leq v_S; \end{aligned}$$

since $\kappa_S \leq \kappa_B$, the solution to this program are the prices given in Theorem 3.

4.2.3 Eliminating Agency Fees

Whereas agent specialization increases the scope for collusion, eliminating agency fees may reduce the scope for collusion. In response to eliminating agency fees, there are two candidate equilibria to maximize industry profits. In the first, seller-proficient agents represent both buyers and sellers. In the second, buyer-exclusive agents represent buyers while seller-proficient agents represent sellers. The advantage of having seller-proficient agents represent buyers is that, in equilibrium, buyers can still be charged negative prices as seller-proficient agents' profits from representing sellers will more than recover the costs of representing buyers. Note that buyer-exclusive agents cannot profitably undercut, as it is easy to incentivize seller-proficient agents to ostracize them.

Theorem 4. *For $\delta \geq \frac{1}{2}$, the highest sustainable industry profits (and the highest sustainable profits for seller-proficient agents) when buyers only work with seller-proficient agents are achieved with prices*

$$p_B^* = \begin{cases} v_B & \hat{\sigma} \geq \kappa_S \\ \frac{v_B - \kappa_S(v_B + v_S - c) + \hat{\sigma}(v_S - c)}{1 - \hat{\sigma}} & \hat{\sigma} \leq \kappa_S \end{cases}$$

$$p_S^* = v_S$$

where $\hat{\sigma} = \frac{\sigma}{(1-\delta)\kappa_B}$.

Moreover, industry revenue is weakly less than the highest sustainable profits when agency fees are allowed. However, $\lim_{\sigma \rightarrow 0}(p_B^* + p_S^*) = (v_B + v_S)(1 - \kappa_S) + \kappa_S c$, which is the same revenue in the limit as under the equilibrium which supports the highest sustainable profits when agency fees are allowed.

The equilibrium supporting the prices stated in Theorem 4 is similar to the equilibrium supporting the prices when agent specialization is not present, i.e., the equilibrium described after Theorem 1. However, in the equilibrium described by Theorem 4, there is an additional incentive constraint: We may need to incentivize seller-proficient agents to represent buyers each period through dynamic considerations, because when p_B^* is less than c (i.e., the price charged to buyers does not compensate the agent for the cost of representing that buyer), a seller-proficient agent would be better off not representing buyers. The additional constraint does not bind so long as $\delta \geq \frac{1}{2}$.

As demonstrated in Figure 5, the equilibrium price is lower under the equilibrium of Theorem 4 than under Theorem 3 (in which buyer-exclusive agents represent buyers and are compensated through agency fees). The equilibrium price for buyers is lower as individual equilibrium profits are lower (since seller-proficient agents now inefficiently




Figure 5: The prices supporting the highest sustainable profit equilibrium when seller-proficient agents work with buyers. The dark red line is the seller price, p_S^* , and the blue line is the buyer price, p_B^* . The dashed green line is the buyer price when agency fees are allowed; the seller price is identical in both cases. Here, $\delta = \frac{3}{4}$, $v_B = 1$, $v_S = 10$, $\kappa_B = \frac{1}{4}$, $\kappa_S = \frac{1}{5}$, and $c = 1$.

represent buyers). But individual equilibrium profits go to 0 as the market becomes highly unconcentrated regardless of whether agency fees are allowed; thus, as the market becomes highly unconcentrated, the highest sustainable buyer price when agency fees are allowed (given in Theorem 3) and the highest sustainable buyer price when buyers work with seller-proficient agents (given in Theorem 4) converge.

Theorem 5. *For $\delta \geq \frac{1}{2}$, the highest sustainable industry profits (and the highest sustainable profits for seller-proficient agents) when buyers work with buyer-exclusive agents are achieved with prices*

$$p_B^* = \begin{cases} v_B & \hat{\sigma} \geq \kappa_S \left(1 - \frac{c}{v_B + v_S}\right) \\ \frac{v_B - \kappa_S(v_B + v_S - c) + \hat{\sigma}v_S}{1 - \hat{\sigma}} & \hat{\sigma} \in \left[\kappa_S \left(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1 - \kappa_S}{\kappa_S}\right), \kappa_S \left(1 - \frac{c}{v_B + v_S}\right)\right] \\ 0 & \hat{\sigma} \leq \kappa_S \left(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1 - \kappa_S}{\kappa_S}\right) \end{cases}$$

$$p_S^* = \begin{cases} v_S & \hat{\sigma} \geq \kappa_S \left(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1 - \kappa_S}{\kappa_S}\right) \\ \frac{v_B(1 - \kappa_S) + \frac{\kappa_S}{\kappa_B} v_S(1 - \kappa_B) + \kappa_S c}{\frac{\kappa_S}{\kappa_B} - \hat{\sigma}} & \hat{\sigma} \leq \kappa_S \left(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1 - \kappa_S}{\kappa_S}\right) \end{cases}$$

where $\hat{\sigma} = \frac{\sigma}{(1 - \delta)\kappa_B}$.²⁸

²⁸Note that we can only have $\hat{\sigma} = \frac{\sigma}{(1 - \delta)\kappa_B} \leq \kappa_S \left(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1 - \kappa_S}{\kappa_S}\right)$ if the buyers' valuations are sufficiently low, i.e., $v_B \leq \frac{\kappa_S}{1 - \kappa_S}(v_S - c)$.




Figure 6: The prices supporting the highest sustainable profit equilibrium when seller-proficient agents work with buyers. The orange line is the seller price, p_S^* , and the blue line is the buyer price, p_B^* . The dashed green line is the buyer price when agency fees are allowed; the dashed red line is the seller price when agency fees are allowed. Here, $\delta = \frac{3}{4}$, $v_B = 1$, $v_S = 10$, $\kappa_B = \frac{1}{4}$, $\kappa_S = \frac{1}{5}$, and $c = 1$.

Moreover, industry revenue is weakly less than in the highest sustainable profit equilibrium with agency fees.

The interesting case of Theorem 5 is when buyer valuations are sufficiently low; in that case, buyers are charged a 0 price and sellers are charged a price *less than* v_S . Figure 6 exhibits this effect: Without agency fees, buyers are charged a 0 price as it is the lowest price for which buyer-exclusive agents will represent buyers. This hard floor for buyer prices limits the set of sustainable seller prices: for any seller price higher than p_S^* , a seller-proficient agent could choose prices sufficiently low to attract buyers and sellers while still increasing profits. Thus, eliminating agency fees may appreciably decrease the scope for collusion, as suggested by Barwick (2018).²⁹

Corollary 1. *If $\delta \geq \frac{1}{2}$ and the value to buyers relative to sellers is sufficiently low, i.e., $v_B \frac{1-\kappa_S}{\kappa_S} + c < v_S$, then the prices described in Theorem 5 for an unconcentrated industry— $\lim_{\sigma \rightarrow 0} (p_B^* + p_S^*) = v_B \left(\frac{\kappa_B}{\kappa_S} - 1 \right) + v_S (1 - \kappa_B) + \kappa_B c$ —are less than the prices under the highest sustainable profit equilibrium with agency fees for an unconcentrated industry (Theorem 3).*

²⁹Note that whether the equilibrium of Theorem 4 or the equilibrium of Theorem 5 is more profitable depends on the cost for seller-proficient agents of representing buyers. It is straightforward to construct examples where c is sufficiently large that all agents prefer the equilibrium of Theorem 5 when agency fees are eliminated.

4.3 Encouraging For Sale by Owner Properties and Buyer Self-Representation

An especially simple proposed solution is to simply encourage more “For Sale by Owner” properties; similarly, some have proposed that buyers represent themselves. However, neither of these solutions has seen widespread adoption,³⁰ and the dynamics we model here can help us understand why.

If we allowed buyers to represent themselves in our model, real estate agents can still easily enforce the collusive commission rate by either refusing to work with buyers who self-represent or demanding the entire commission in transactions with those buyers. In fact, in many states it is illegal for non-licensed agents to collect a commission (even when representing themselves).³¹

It is similarly difficult for homeowners to represent themselves in real estate transactions. Should a homeowner do so, buyer agents can ostracize that homeowner just as they would ostracize a price-deviating seller agent; indeed, consistent with this prediction, for-sale-by-owner homes take much longer to sell than sales via agents (Hendel et al., 2009). And, since most sellers are simultaneously buying a house (and are unwilling or unable to pay two mortgages simultaneously), such delay is very costly. Indeed, arm’s-length sales of for-sale-by-owner homes make up only 5% of the market and are in decline (hitting an all-time low in 2017) (National Association of Realtors, 2017).

5 Conclusion

Our analysis explains how an extremely unconcentrated industry such as real estate brokerage can still support collusive pricing. The brokerage feature is key: our analysis relies on the fact that, after prices have been announced, brokers must work with each other to complete transactions. Brokers can punish a price-deviator within-period by refusing to work with him; this both directly harms the price-deviator and makes the price-deviator’s services less appealing to buyers and sellers. As a result, even though the market has low barriers to entry (and thus many brokers), brokers are able to extract a large fraction of total surplus.

Our results suggest that there is substantial potential for intermediaries in residential real estate to extract rents; this rent extraction comes in the form of a high price for brokering transactions. Given that residential real estate is a key component of the U.S. economy, with approximately one and a half trillion dollars in transactions per year

³⁰According to the National Association of Realtors (2017), only 5% of real estate transactions were sold directly by the owner to a non-personal acquaintance.

³¹See, e.g., Florida statute §475.25(1)(h), Texas statute §1101.651, and New York Real Property Law §442.

([National Association of Realtors, 2018](#)), if housing demand is at least somewhat elastic, the allocative distortions resulting from broker rents could be quite economically significant.³²

By formally modeling brokered intermediation, we can also clarify the roles of proposed and existing policies. Our model indicates that eliminating rebate bans, while not a panacea, reduces the scope for collusion. Similarly, eliminating agency fees can also weaken the ability of industry participants to maintain high prices.

Finally, our work highlights the importance of using repeated extensive form games to model competition in settings with multi-stage interactions among market participants.³³ The strategies analyzed here do not fit the traditional normal form repeated game analysis paradigm, which may partially explain why many economists consider the high commissions of real estate agents puzzling ([Hsieh and Moretti, 2003](#)). Moreover, multi-stage interactions are quite common in intermediated markets, including venture capital, initial public offerings, and syndicated lending ([Hochberg et al., 2010](#); [Hatfield et al., 2020](#); [Cai et al., 2020](#)). These techniques may also prove useful in analyzing business-to-business transactions and market entry.

³²See [Barwick and Wong \(2019\)](#) for an argument that the high commission rates in the United States have significant welfare consequences.

³³Repeated extensive form games have also been used to analyze vertical mergers ([Nocke and White, 2007](#)) and markets with syndicated production ([Hatfield et al., 2020, 2021](#)).

References

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica* 54(4), 383–396. (Cited on page 4.)
- Ali, S. N. and C. Liu (2019). Conventions and coalitions in repeated games. arXiv:1906.00280. (Cited on page 10.)
- Ambrus, A. (2006). Coalitional rationalizability. *Quarterly Journal of Economics* 121(3), 903–929. (Cited on page 10.)
- Athey, S. and K. Bagwell (2001). Optimal collusion with private information. *The RAND Journal of Economics* 32(3), 428–465. (Cited on page 3.)
- Athey, S., K. Bagwell, and C. Sanchirico (2004). Collusion and price rigidity. *Review of Economic Studies* 71(2), 317–349. (Cited on page 3.)
- Barwick, P. (2018). Opening remarks on regulatory and industry factors affecting residential real estate competition. What’s New in Residential Real Estate Brokerage Competition – An FTC-DOJ Workshop (<https://www.ftc.gov/news-events/events-calendar/2018/04/whats-new-residential-real-estate-brokerage-competition-ftc-doj>). (Cited on pages 5, 26, and 34.)
- Barwick, P. J. and P. A. Pathak (2015). The costs of free entry: An empirical study of real estate agents in Greater Boston. *The RAND Journal of Economics* 46(1), 103–145. (Cited on page 2.)
- Barwick, P. J., P. A. Pathak, and M. Wong (2017). Conflicts of interest and steering in residential brokerage. *American Economic Journal: Applied Economics* 9(3), 191–222. (Cited on pages 2 and 4.)
- Barwick, P. J. and M. Wong (2019). Competition in the real estate brokerage industry: A critical review. Working paper, Cornell University. (Cited on page 36.)
- Beck, J., F. Scott, and A. Yelowitz (2012). Concentration and market structure in local real estate markets. *Real Estate Economics* 40(3), 422–460. (Cited on page 2.)
- Bernheim, B. D. and J. Meer (2013). Do real estate brokers add value when listing services are unbundled? *Economic inquiry* 51(2), 1166–1182. (Cited on page 2.)
- Bernheim, B. D., B. Peleg, and M. D. Whinston (1987). Coalition-proof Nash equilibria I. concepts. *Journal of Economic Theory* 42(1), 1–12. (Cited on page 9.)

- Birger, J. and J. Caplin (2004). The 4 1/2% solution. *Money* 33(10), 104–109. (Cited on page 3.)
- Cai, J., F. Eidam, A. Saunders, and S. Steffen (2020). Loan syndication structures and price collusion. SSRN #3250817. (Cited on page 36.)
- Cestau, D., B. Hollifield, D. Li, and N. Schürhoff (2019). Municipal bond markets. *Annual Review of Financial Economics* 11(1), 65–84. (Cited on page 5.)
- Christie, W. G. and P. H. Schultz (1994). Why do nasdaq market makers avoid odd-eighth quotes? *The Journal of Finance* 49(5), 1813–1840. (Cited on page 5.)
- Christie, W. G. and R. B. Thompson (2006). Wall street scandals: The curative effects of law and finance. *Washington University Law Review* 84, 1567–1590. (Cited on page 5.)
- Delcoure, N. and N. G. Miller (2002). International residential real estate brokerage fees and implications for the us brokerage industry. *International Real Estate Review* 5(1), 12–39. (Cited on pages 2 and 21.)
- Department of Justice (2005). Complaint. U.S. v. Kentucky Real Estate Commission. (Cited on pages 4 and 41.)
- Department of Justice and Federal Trade Commission (2007). Competition in the real estate brokerage industry. <https://www.justice.gov/sites/default/files/atr/legacy/2007/05/08/223094.pdf>. (Cited on pages 3 and 4.)
- Durr, K. and R. Colby (2010). The institution of experience: Self-regulatory organizations in the securities industry, 1972-2010. <http://www.sechistorical.org/museum/galleries/sro/index.php>. (Cited on page 5.)
- Federal Trade Commission (1983). *The residential real estate brokerage industry : a staff report*. Federal Trade Commission. (Cited on page 3.)
- Hall, R. E. and P. R. Milgrom (2008). The limited influence of unemployment on the wage bargain. *American economic review* 98(4), 1653–74. (Cited on page 8.)
- Han, L. and S.-H. Hong (2011). Testing cost inefficiency under free entry in the real estate brokerage industry. *Journal of Business & Economic Statistics* 29(4), 564–578. (Cited on pages 4 and 24.)

- Hatfield, J. W., S. D. Kominers, and R. Lowery (In preparation, 2021). Collusion with hidden syndication recruiting. University of Texas working paper. (Cited on page 36.)
- Hatfield, J. W., S. D. Kominers, R. Lowery, and J. M. Barry (Forthcoming, 2020). Collusion in markets with syndication. *Journal of Political Economy*. (Cited on page 36.)
- Hendel, I., A. Nevo, and F. Ortalo-Magné (2009). The relative performance of real estate marketing platforms: MLS versus FSBOMadison.com. *American Economic Review* 99(5), 1878–98. (Cited on page 35.)
- Hochberg, Y. V., A. Ljungqvist, and Y. Lu (2010). Networking as a barrier to entry and the competitive supply of venture capital. *Journal of Finance* 65(3), 829–859. (Cited on page 36.)
- Hsieh, C.-T. and E. Moretti (2003). Can free entry be inefficient? Fixed commissions and social waste in the real estate industry. *Journal of Political Economy* 111(5), 1076–1122. (Cited on pages 2 and 36.)
- Kang, A. and R. Lowery (2014). The pricing of IPO services and issues: Theory and estimation. *Review of Corporate Finance Studies* 2(2), 188–234. (Cited on page 3.)
- Levitt, S. D. and C. Syverson (2008a). Antitrust implications of outcomes when home sellers use flat-fee real estate agents. *Brookings-Wharton Papers on Urban Affairs*, 47–93. (Cited on page 4.)
- Levitt, S. D. and C. Syverson (2008b). Market distortions when agents are better informed: The value of information in real estate transactions. *Review of Economics and Statistics* 90(4), 599–611. (Cited on page 2.)
- Magura, M. (2007). How rebate bans, discriminatory MLS listing policies, and minimum service requirements can reduce competition for real estate brokerage services and why it matters. Economic Analysis Group Discussion Paper (EAG07-8), Antitrust Division, Department of Justice. (Cited on page 24.)
- Mailath, G. J., V. Nocke, and L. White (2017). When and how the punishment must fit the crime. *International Economic Review* 58(2), 315–330. (Cited on pages 4 and 12.)
- National Association of Realtors (2017). 2017 profile of home buyers and sellers. (Cited on page 35.)

- National Association of Realtors (2018). Existing-home sales overview. <https://www.nar.realtor/sites/default/files/documents/ehs-09-2018-overview-2018-10-19.pdf>. (Cited on page 36.)
- Nocke, V. and L. White (2007). Do vertical mergers facilitate upstream collusion? *American Economic Review* 97(4), 1321–1339. (Cited on page 36.)
- Smith, A. (1776). *An inquiry into the nature and causes of the Wealth of Nations*. W. Strahan and T. Cadell. (Cited on page 2.)
- Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review* 100(4), 1642–72. (Cited on page 9.)
- White, A. and E. G. Weyl (2016). Insulated platform competition. Working paper, Tsinghua University. (Cited on page 9.)

A Candid Remarks by Real Estate Agents

All remarks by real estate agents provided here are as documented by the [Department of Justice \(2005\)](#) in their complaint in *U.S.A. v. Kentucky Real Estate Commission*. Many real estate agents argued that the Kentucky rebate ban should not be eliminated, expressly because it forestalled lower prices:

- “If we give rebates and inducements, it would get out of control and all clients would be wanting something. The present law keeps it under control.”
- “I am for the law as it stands now. If inducements were allowed, they could lead to competitive behavior, which would make us look unprofessional in the eyes of the public.”
- “I think this would just take money right out of our pocket.”

Moreover, many agents recognized the explicitly anti-competitive nature of the ban:

- “Rebates and inducements will increase competition and give consumers more choices in service.”
- “Current law inhibits free trade.”
- “Commissions and sales awards are common in other industries. The bigger wrong being committed by agents and brokers is the informal unspoken price fixing that occurs.”

B Random Rationing

Here, we define an algorithm for allocating buyers to agents given a preference list for each buyer. The procedure for sellers is analogous.

We define the assignment of buyers to agents recursively as

$$\mathbf{a}(b) = \max_{\succ_b} \left(\{a \in A : \mathbf{a}^{-1}([0, b]) \leq \kappa_B\} \cup \{\emptyset\} \right).$$

where $a \succ_b \hat{a}$ if a is ranked higher than \hat{a} in b 's preference list (and the outside option \emptyset is listed immediately after all acceptable agents).

C Proofs

C.1 Proof of Proposition 1

Existence of a Zero-Profit Equilibrium of the Stage Game

We define the *Bertrand competition strategy profile* under which every agent obtains a payoff of 0:

1. Each agent a offers a buyer price $p_B^a = 0$ and a seller price $p_S^a = 0$.
2. If every agent offers non-negative buyer and seller prices, then:
 - Each buyer reports a ranking over all agents a who offer a buyer price $p_B^a \in [0, v_B]$, where agents offering lower buyer prices are listed before agents offering higher buyer prices; each such ranking is reported with equal probability by each buyer.
 - Each seller reports a ranking over all agents a who offer a seller price $p_S^a \in [0, v_S]$, where agents offering lower seller prices are listed before agents offering higher seller prices; each such ranking is reported with equal probability by each seller.

Otherwise, if some agent offers a negative buyer or seller price, buyers and sellers play coordination-proof Nash equilibrium strategies of the subgame in Step 2 (given the network formation in Steps 3 and 4 of this period).³⁴

3. Regardless of the price offers, each agent a invites every other agent \tilde{a} , offering a transfer of $f^{\tilde{a} \leftarrow a} = -\max\{p_B^{\tilde{a}}, -p_S^a\}$; that is, a demands either all of the surplus that \tilde{a} will receive from a buyer (i.e., $-p_B^{\tilde{a}}$) or surplus sufficient to ensure that a 's profits from a transaction are non-negative.
4. Regardless of the price offers, each agent \tilde{a} accepts the invitation from a so long as $f^{\tilde{a} \leftarrow a} \geq -p_B^{\tilde{a}}$.

Agents' actions with respect to accepting or rejecting invitations are clearly optimal, as an agent accepts an invitation if and only if any transaction facilitated by that invitation provides that agent with non-negative surplus. It is then immediate that agents' actions with respect to making invitations are optimal, as each invitation made by a to \tilde{a} is either the lowest fee offer that will be accepted by \tilde{a} or is an offer that will be rejected by \tilde{a} (and no fee that obtains a positive payoff for a will be accepted).

³⁴Note that here and in the sequel, players have correct conjectures about network formation via backward induction.

It is also clear that buyer and seller actions are optimal and coordination-proof; when no agent offers a negative price, it is immediate that the full network forms. Thus, buyers and sellers prefer lower prices, as every agent offers access to the entire other side of the market. If any agent offers a negative buyer or seller price, then we simply require that buyers and sellers play coordination-proof Nash equilibrium strategies.

Finally, agents' price offers are optimal; given that every other agent offers a buyer price of 0 and a seller price of 0, if a offers a positive buyer price, a will not represent any buyers and, if a offers a positive seller price, a will not represent any sellers.³⁵ Thus, a cannot increase his profits by increasing his buyer or seller price. It is immediate that a cannot increase his profits by decreasing his buyer or seller price.³⁶

Proof that Profits Must Be 0 in the Stage Game

We now show that there does not exist any positive-profit symmetric subgame-perfect equilibrium of the stage game.

Consider a positive profit symmetric subgame-perfect equilibrium of the stage game; we will show that any such equilibrium is not coordination-proof. In any such equilibrium, each agent $a \in A$ is offering the same price p_B to buyers and the same price p_S to sellers; moreover $p_B + p_S > 0$.

Lemma 1. *Consider the network formation subgame. If $p_B^a + p_S^{\bar{a}} > 0$, $\mathbf{B}(a) > 0$, and $\mathbf{S}(\bar{a}) > 0$, then $a \in A^{\Rightarrow \bar{a}}$. Moreover, $f^{a \leftarrow \bar{a}} = -p_B^a$.*

Proof. Agent a will accept any fee $f^{a \leftarrow \bar{a}} > -p_B^a$ as this increases the payoff for a (since $\mathbf{B}(a) > 0$ and $\mathbf{S}(\bar{a}) > 0$). Thus, if \bar{a} offers a fee such that a rejects \bar{a} 's invitation, \bar{a} can offer $-p_B^a + \epsilon$ which a will accept if $\epsilon > 0$. But, for ϵ small enough, offering such a fee and having it accepted strictly increases \bar{a} 's profits.

If $f^{a \leftarrow \bar{a}} > -p_B^a$, then \bar{a} could increase his profits by choosing to offer a a fee of $f^{a \leftarrow \bar{a}} - \epsilon$ (where $\epsilon > 0$) if this invitation was accepted; agent a would still accept this fee so long as ϵ is sufficiently small. \square

Lemma 1 implies that every agent with a positive mass of buyers accepts the fee of every other agent with a positive mass of sellers since $p_B + p_S > 0$.

There exists at least one agent a such that $\mathbf{S}(a) \equiv \mu_S^a < \kappa_S$. Suppose that a deviates to offer a seller price of $p_S - \epsilon$ (where $\epsilon > 0$). After such a deviation, if $\mathbf{S}(a) > 0$ and ϵ is sufficiently small, Lemma 1 implies that every agent with a positive mass of buyers accepts

³⁵Recall that we have assumed that there is sufficient capacity to serve all of the buyers and sellers even if one agent leaves the economy, i.e., $(|A| - 1)\kappa_B \geq 1$ and $(|A| - 1)\kappa_S \geq 1$.

³⁶We show in Appendix C.1 that all symmetric coordination-proof Nash equilibria have 0 profits.

an invitation with a fee of $-p_B$ from a . It is then immediate that a mass of sellers of size κ_S is strictly better off by listing a first, since those sellers receive a better price and still have access to all the buyers. Thus, the profits of agent a have increased as (for ϵ sufficiently small)

$$\mu_S^a(p_B + p_S) < \kappa_S(p_B + p_S - \epsilon).$$

C.2 Details of the Proof of Theorem 1

Making and Responding to Invitations in the Cooperation Phase

There are four cases to consider:

Case 1: Collusive Pricing. For an agent a , if either

- a offered a fee other than $-p_B^*$ (or did not invite some agent), or
- some agent offered a a fee other than $-p_B^*$,

then play will proceed to the Bertrand reversion phase (regardless of a 's actions at this point); thus, it is immediate that a will accept an invitation if and only if the fee is at least $-p_B^*$ (as this maximizes in-period profits). Otherwise, it is optimal for a to accept each offered fee of $-p_B^*$ as this has no effect on in-period profits and, by doing so, ensures that play continues in the cooperation phase.

It is also immediate that it is optimal to invite every other agent with a fee of $-p_B^*$ as this maximizes in-period profits as well as ensuring that play continues in the cooperation phase.

Case 2: Ineffective price deviation by \hat{a} . Showing that all agents other than \hat{a} follow prescribed play with respect to making invitations to and accepting invitations from agents other than \hat{a} is exactly as in the collusive pricing case. It is also immediate that a should follow the prescribed strategy with respect to any invitation from \hat{a} , as \hat{a} has no sellers (so this period's profits are unaffected). Furthermore, a should follow the prescribed strategy of not making an invitation to \hat{a} , as \hat{a} has no buyers (so this period's profits are unaffected). Finally, any set of invitations and fees is optimal for \hat{a} since \hat{a} has no sellers and receives 0 following equilibrium continuation play; moreover, any acceptance/rejection of invitations is optimal for \hat{a} since \hat{a} has no buyers and receives 0 following equilibrium continuation play.

Case 3: Effective price deviation by \hat{a} . Showing that all agents other than \hat{a} follow prescribed play with respect to making invitations to and accepting invitations from

agents other than \hat{a} is exactly as in the collusive pricing case. Moreover, it is immediate that \hat{a} 's prescribed actions are optimal:

- Making any invitation with a fee greater than $p_S^{\hat{a}}$ is not optimal as such an invitation will be accepted but result in lower profits this period and have no effect on profits in future periods for \hat{a} .
- Making any invitation with a fee less than $p_S^{\hat{a}}$ is not optimal as no such invitation will be accepted and have no effect on profits in future periods for \hat{a} .
- Accepting an invitation with a fee less than $-p_B^{\hat{a}}$ is not optimal as it results in lower profits this period and has no effect on profits in future periods.
- Not accepting an invitation with a fee greater than $-p_B^{\hat{a}}$ is not optimal as it results in lower profits this period and has no effect on profits in future periods.

That an agent a is better off following his prescribed actions than if he invited \hat{a} with a fee of $-p_B^{\hat{a}}$ and accepted an invitation from \hat{a} with a fee of $p_S^{\hat{a}}$ (and followed his prescribed actions with respect to other agents) is proven in the text.

Making and Responding to Invitations in the Collusive Punishment Phase

Cases 1 and 2: Collusive Pricing and Ineffective Price Deviations by \hat{a} .

The analysis here is analogous to that of Case 2 during the cooperation phase.

Case 3: Effective price deviation by \hat{a} . The analysis here follows as in the analysis of Case 3 of the cooperation phase, except that the in-period profits from working with other agents now depend on q_B^* and q_S^* instead of p_B^* and p_S^* . In particular, the total payoff for a from following his prescribed actions is (cf. (2))

$$\underbrace{\underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)}_{\text{Mass of sellers represented by } a} \underbrace{(1-\kappa_B)}_{\text{Mass of buyers represented by agents other than } \hat{a}} \underbrace{(q_B^* + q_S^*)}_{\text{Profit per transaction}}}_{\text{Profits from working with agents other than } \hat{a} \text{ this period}} + \underbrace{\frac{\delta}{1-\delta} \frac{\alpha}{1-\alpha} (q_B^* + q_S^*)}_{\text{Payoff in future periods from adhering}}$$

while the payoff for working with \hat{a} is given by (cf. (1))

$$\begin{aligned}
& \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)(1-\kappa_B)(q_B^* + q_S^*)}_{\text{Profits from working with agents other than } \hat{a} \text{ this period}} + \\
& \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)}_{\text{Mass of sellers represented by } a} \underbrace{\kappa_B}_{\text{Mass of buyers represented by } \hat{a}} \underbrace{(p_B^{\hat{a}} + q_S^*)}_{\text{Profit per transaction}} + \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_B)}_{\text{Mass of buyers represented by } a} \underbrace{\kappa_S}_{\text{Mass of sellers represented by } \hat{a}} \underbrace{(q_B^* + p_S^{\hat{a}})}_{\text{Profit per transaction}}. \\
& \hspace{10em} \underbrace{\hspace{10em}}_{\text{Payoff from working with } \hat{a}}
\end{aligned}$$

Since $p_B^{\hat{a}} \leq q_B^*$ and $p_S^{\hat{a}} \leq q_S^*$, the analysis that it is optimal for a to reject working with \hat{a} so long as $\delta \geq \frac{1}{2}$ is analogous.

Agent Selection in the Collusive Punishment Phase

The analysis after collusive pricing (Case 1) is identical. After a effective or ineffective price deviation by an agent \hat{a} (Cases 2–3) the analysis is similar to that of price deviations in the cooperation phase, except that the other agents are now offering prices of q_B^* and q_S^* instead of p_B^* and p_S^* . Thus, we find that, for a price deviation to be effective during a collusive punishment phase, we need that

$$p_B^{\hat{a}} \leq v_B - \frac{1}{\kappa_S}(v_B - q_B^*) = q_B^\circ \text{ and } p_S^{\hat{a}} \leq v_S - \frac{1}{\kappa_B}(v_S - q_S^*) = q_S^\circ.$$

C.3 Proof of Theorem 2

The proof proceeds as the proof of Theorem 1 except that the buyer deviation prices are now

$$\begin{aligned}
p_B^\circ &= \max\left\{0, v_B - \frac{1}{\kappa_S}(v_B - p_B^*)\right\} \\
q_B^\circ &= \max\left\{0, v_B - \frac{1}{\kappa_S}(v_B - q_B^*)\right\}.
\end{aligned}$$

Intuitively, these buyer deviation prices reflect the fact that, under a rebate ban, no agent (including a price-deviating agent) can charge a price less than 0; thus, the most tempting price offer you can make to buyers after a price deviation is 0. The proof follows *mutatis mutandis* to the proof of Theorem 1 with these substitutions.³⁷

³⁷Note that in our proof that agents will not work with a price deviator during the cooperation phase, we used the fact that $p_B^* + p_S^* \leq v_B + v_S$ and $p_B^{\hat{a}} + p_S^{\hat{a}} \leq v_B + v_S$ to place an upper bound on the profit per transaction; see (1) and (2). A similar argument holds for the collusive punishment phase.

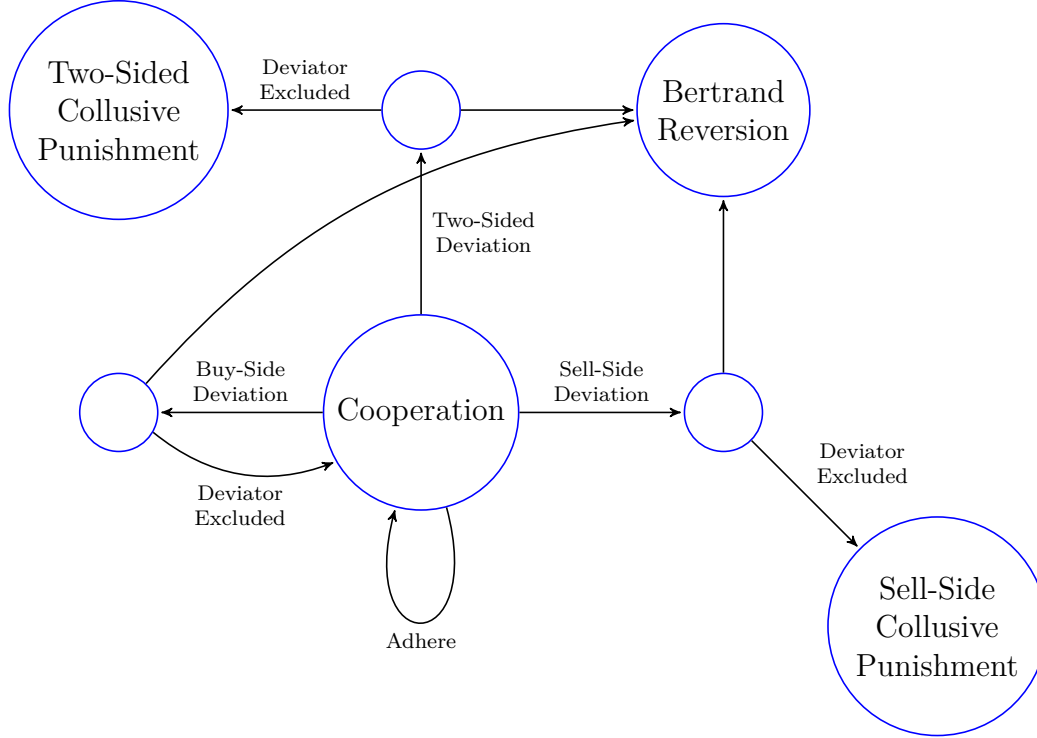


Figure 7: Simplified automaton representation of the equilibrium we consider; in particular, we do not show how play may evolve from the sell-side and two-sided collusive punishment phases. Labeled nodes are phases; unlabeled nodes are intermediate phases, which represent the branching of transitions based on behavior in the later steps of the game.

C.4 Proof of Theorem 3

We now formally construct a strategy profile that sustains (p_B^*, p_S^*) . To simplify the exposition, we first define distinguished actions for buyers and sellers:

1. A buyer *lists buyer-exclusive agents arbitrarily* by reporting with equal probability each ranking that includes every buyer-exclusive agent; similarly, a seller *lists seller-proficient agents arbitrarily* by reporting with equal probability each ranking that includes every seller-proficient agent.
2. A buyer *prioritizes agent a among buyer-exclusive agents* by reporting with equal probability each ranking that both includes every buyer-exclusive agent and ranks a a first; similarly, *prioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a a first.
3. A buyer *deprioritizes agent a among buyer-exclusive agents* by reporting with equal probability each ranking that both includes every buyer-exclusive agent and ranks a

last; similarly, *deprioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a last.

We further define distinguished actions for agents in the network formation steps:

1. The *regular network forms with standard fees* when:
 - Each seller-proficient agent a invites each buyer-exclusive agent \tilde{a} with a fee that demands \tilde{a} 's buyer price when transactions create non-negative surplus, i.e., a fee of $-\max\{p_B^{\tilde{a}}, -p_S^a\}$ and does not invite any other seller-proficient agent.
 - Each seller-proficient agent rejects any invitation from another seller-proficient agent.
 - Each buyer-exclusive agent \tilde{a} accepts every invitation with a fee greater than or equal to $-p_B^{\tilde{a}}$.

2. The *network excluding (the seller-proficient agent) a forms with standard fees* when:
 - Each seller-proficient agent other than a invites each buyer-exclusive agent \tilde{a} with a fee that demands \tilde{a} 's buyer price when transactions create non-negative surplus, i.e., a fee of $-\max\{p_B^{\tilde{a}}, -p_S^a\}$, and every buyer-exclusive agent \tilde{a} accepts every invitation with an offer greater than or equal to $-p_B^{\tilde{a}}$. Moreover, each seller-proficient agent other than a does not invite any other seller-proficient agent and each seller-proficient agent other than a rejects any invitation from another seller-proficient agent.
 - The seller-proficient agent a invites each buyer-exclusive agent with a fee of p_S^a and does not invite any other seller-proficient agent. The seller-proficient agent a accepts the invitation from any other seller-proficient agent \check{a} if and only if $p_B^a + f^{a \leftarrow \check{a}} - c \geq 0$.
 - Each buyer-exclusive agent \tilde{a} accepts every invitation from a seller-proficient agent other than a with a fee greater than or equal to $-p_B^{\tilde{a}}$; \tilde{a} accepts an invitation from a if and only if $f^{\tilde{a} \leftarrow a} > p_S^a$.

Thus, the regular network forms among agents other than a and no agent forms any links with a .

3. The *network treating (the seller-proficient agent) a as a buyers' agent forms with standard fees* when:

- Each seller-proficient agent \hat{a} other than a invites each buyer-exclusive agent \tilde{a} with a fee that demands \tilde{a} 's buyer price when transactions create non-negative surplus, i.e., a fee of $-\max\{p_B^{\tilde{a}}, -p_S^{\hat{a}}\}$, invites a with a fee of $-\max\{p_B^{\tilde{a}} - c, -p_S^{\hat{a}}\}$, and does not invite any other seller-proficient agent.
- The seller-proficient agent a does not make any invitations.
- Each seller-proficient agent other than a rejects any invitation from another seller-proficient agent.
- The seller-proficient agent a accepts every invitation with a fee greater than or equal to $-p_B^a + c$.
- Each buyer-exclusive agent \tilde{a} accepts every invitation with a fee greater than or equal to $-p_B^{\tilde{a}}$.

We also define the buyer deviation price in the cooperation phase as $p_B^o = v_B - \frac{1}{\kappa_S}(v_B - p_B^*)$ and the seller deviation price $p_S^o = v_S - \frac{1}{\kappa_B}(v_S - p_S^*) = v_S$; these are the prices at which buyers and sellers will be willing to work with a deviating agent in the cooperation phase. During the collusive punishment phase, prices are $q_B^* = (1 - \kappa_S)v_B - \kappa_S v_S + \kappa_S c$ and $q_S^* = v_S$. We define the buyer deviation price during the collusive punishment phase as $q_B^o = v_B - \frac{1}{\kappa_S}(v_B - q_B^*)$ and the seller deviation price as $q_S^o = v_S - \frac{1}{\kappa_B}(v_S - q_S^*) = v_S$; these are the prices at which buyers and sellers will be willing to work with an agent who deviates in the collusive punishment phase. We also define the fees g^* and h^* during the two-sided and sell-side collusive punishment phases, respectively, as in the text. Finally, we suppress here detailing the strategies after mutual deviations, i.e., after two or more agents *simultaneously* deviate: since no agent expects any other agent to deviate, such cases have no effect on incentives.

The strategy profile that sustains p_B^* and p_S^* consists of four phases. In the *cooperation phase*:

1. Each buyer-exclusive agent offers a buyer price p_B^* and each seller-proficient agent offers a seller price p_S^* .
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing: Each buyer-exclusive agent has offered a buyer price p_B^* and each seller-proficient agent has offered a seller price p_S^* . Buyers list buyer-exclusive agents arbitrarily and sellers list seller-proficient agents arbitrarily.

Case 2a: Appealing price deviation by $\check{a} \in A_B$: Each buyer-exclusive agent except \check{a} has offered a buyer price p_B^* , the buyer-exclusive agent \check{a} has offered a buyer price $p_B^{\check{a}} \in [-p_S^*, p_B^*)$, and each seller-proficient agent has offered a seller price p_S^* .

Buyers prioritize \check{a} among buyer-exclusive agents³⁸ and sellers list seller-proficient agents arbitrarily.

Case 2b: *Unappealing price deviation by $\check{a} \in A_B$* : Each buyer-exclusive agent except \check{a} has offered a buyer price p_B^* , the buyer-exclusive agent \check{a} has offered a buyer price $p_B^{\check{a}} \in (-\infty, -p_S^*) \cup (p_B^*, \infty)$,³⁹ and each seller-proficient agent has offered a seller price p_S^* . Buyers deprioritize \check{a} among buyer-exclusive agents and sellers list seller-proficient agents arbitrarily.

Case 3: *Ineffective price deviation by $\hat{a} \in A_S$* : Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent except \hat{a} has offered a seller price p_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \geq (p_B^{\circ}, p_S^{\circ})$ but not $(p_B^{\circ}, p_S^{\circ})$.⁴⁰ Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.

Case 4a: *Appealing buy-side price deviation by $\hat{a} \in A_S$* : Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent except \hat{a} has offered a seller price p_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $p_B^{\hat{a}} \in [-p_S^* + c, p_B^*)$ and $p_S^{\hat{a}} > p_S^{\circ} = v_S$. Buyers prioritize \hat{a} among buyer-exclusive agents and sellers deprioritize \hat{a} among seller-proficient agents.

Case 4b: *Unappealing buy-side price deviation by $\hat{a} \in A_S$* : Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent except \hat{a} has offered a seller price p_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $p_B^{\hat{a}} \in (-\infty, -p_S^* + c) \cup (p_B^*, \infty)$ and $p_S^{\hat{a}} > p_S^{\circ} = v_S$. Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.

Case 5: *Sell-side price deviation by $\hat{a} \in A_S$* : Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent except \hat{a} has offered a seller price p_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $p_B^{\hat{a}} > p_B^{\circ}$ and $p_S^{\hat{a}} < p_S^{\circ} = v_S$. Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.

Case 6: *Effective price deviation by \hat{a}* : Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent except \hat{a} has offered a seller price p_S^* , and

³⁸We could also construct an equilibrium (with the same prices) in which a buyer-exclusive agent who deviates on price does not attract any buyers and is excluded from the network. This requires checking additional incentive constraints and so, for simplicity, we use the strategies delineated here.

³⁹Note that a price deviation can be unappealing for two reasons: If the buyer price is too low, \check{a} will not receive any acceptable invitations from seller-proficient agents. If the buyer price is too high, buyers are better off working with agents offering the lower price p_B^* .

⁴⁰Throughout, when considering off-path price offers, we consider an agent who does not make an offer to buyers (sellers) to have offered buyers (sellers) an infinite buyer (seller) price.

\hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (p_B^o, p_S^o)$. Buyers prioritize \hat{a} among buyer-exclusive agents and sellers prioritize \hat{a} among seller-proficient agents.

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing. The regular network forms with standard fees.

Case 2: Price deviation by a buyer-exclusive agent $\check{a} \in A_B$. The regular network forms with standard fees.⁴¹

Case 3: Ineffective price deviation by $\hat{a} \in A_S$. The network excluding \hat{a} forms with standard fees.

Case 4: Buy-side price deviation by $\hat{a} \in A_S$. The network treating \hat{a} as a buyers' agent forms with standard fees.

Case 5: Sell-side price deviation by $\hat{a} \in A_S$. The network excluding \hat{a} forms with standard fees.

Case 6: Effective price deviation by $\hat{a} \in A_S$. The network excluding \hat{a} forms with standard fees.

4. Under collusive pricing or after a price deviation by a buyer-exclusive agent, if the regular network with standard fees forms, play continues in the cooperation phase. After an effective or ineffective price deviation by $\hat{a} \in A_S$, if the network excluding \hat{a} forms with standard fees, play proceeds to the two-sided \hat{a} -collusive punishment phase. After a buy-side price deviation by $\hat{a} \in A_S$, if the network treating \hat{a} as a buyers' agent forms with standard fees forms, then play continues in the cooperation phase. After a sell-side price deviation by $\hat{a} \in A_S$, if the network excluding \hat{a} forms with standard fees, play proceeds to the sell-side \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the *two-sided \hat{a} -collusive punishment phase*:

1. Each buyer-exclusive agent offers a buyer price q_B^* and each seller-proficient agent offers a seller price q_S^* .
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

⁴¹Recall that both “standard fees” as well as the “regular network” itself depend on the prices offered by buyer-exclusive agents.

- Case 1: *Collusive pricing*:** Each buyer-exclusive agent has offered a buyer price q_B^* and each seller-proficient agent has offered a seller price q_S^* . Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 2a: *Appealing price deviation by $\check{a} \in A_B$* :** Each buyer-exclusive agent except \check{a} has offered a buyer price q_B^* , the buyer-exclusive agent \check{a} has offered a buyer price $p_B^{\check{a}} \in [-q_S^*, q_B^*)$, and each seller-proficient agent has offered a seller price q_S^* . Buyers prioritize \check{a} among buyer-exclusive agents and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 2b: *Unappealing price deviation by $\check{a} \in A_B$* :** Each buyer-exclusive agent except \check{a} has offered a buyer price q_B^* , the buyer-exclusive agent \check{a} has offered a buyer price $p_B^{\check{a}} \in (-\infty, -q_S^*) \cup (q_B^*, \infty)$, and each seller-proficient agent has offered a seller price q_S^* . Buyers deprioritize \check{a} among buyer-exclusive agents and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 3: *Ineffective price deviation by $\hat{a} \in A_S$* :** Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent except \hat{a} has offered a seller price q_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \geq (q_B^*, q_S^*)$ but not (q_B^*, q_S^*) . Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 4a: *Appealing buy-side price deviation by $\hat{a} \in A_S$* :** Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent except \hat{a} has offered a seller price q_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $p_B^{\hat{a}} \in [-q_S^* + c, q_B^*)$ and $p_S^{\hat{a}} > q_S^* = v_S$. Buyers prioritize \hat{a} among buyer-exclusive agents and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 4b: *Unappealing buy-side price deviation by $\hat{a} \in A_S$* :** Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent except \hat{a} has offered a seller price q_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $p_B^{\hat{a}} \in (-\infty, -q_S^* + c) \cup (q_B^*, \infty)$ and $p_S^{\hat{a}} > q_S^* = v_S$. Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 5: *Sell-side price deviation by $\hat{a} \in A_S$* :** Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent except \hat{a} has offered a seller price q_S^* , and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $p_B^{\hat{a}} > q_B^*$ and $p_S^{\hat{a}} < q_S^* = v_S$. Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.
- Case 6: *Effective price deviation by \hat{a}* :** Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent except \hat{a} has offered a seller price q_S^* , and

\hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (q_B^o, q_S^o)$. Buyers prioritize \hat{a} among buyer-exclusive agents and sellers prioritize \hat{a} among seller-proficient agents.

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing. The network excluding \hat{a} forms with standard fees.

Case 2: Price deviation by a buyer-exclusive agent $\check{a} \in A_B$. The network excluding \hat{a} forms with standard fees.

Case 3: Ineffective price deviation by $\hat{a} \in A_S$. The network excluding \hat{a} forms with standard fees.

Case 4: Buy-side price deviation by $\hat{a} \in A_S$. The network treating \hat{a} as a buyers' agent forms with standard fees.

Case 5: Sell-side price deviation by $\hat{a} \in A_S$. The network excluding \hat{a} forms with standard fees.

Case 6: Effective price deviation by $\hat{a} \in A_S$. The network excluding \hat{a} forms with standard fees.

4. Under collusive pricing or after a price deviation by a buyer-exclusive agent, if the network (including fees) implied by equilibrium play forms, then play continues in the two-sided \hat{a} -collusive punishment phase. After an effective or ineffective price deviation by $\hat{a} \in A_S$, if the network excluding \hat{a} forms with standard fees forms, then play proceeds to the two-sided \hat{a} -collusive punishment phase. After a buy-side price deviation by $\hat{a} \in A_S$, if network treating \hat{a} as a buyers' agent forms with standard fees, then play continues in the two-sided \hat{a} -collusive punishment phase. After a sell-side price deviation by $\hat{a} \in A_S$, if the network excluding \hat{a} forms with standard fees, then play proceeds to the sell-side \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

The *sell-side \hat{a} -collusive punishment phase* proceeds exactly as the two-sided \hat{a} -collusive punishment phase except that, under collusive pricing, each seller-proficient agent invites each buyer-exclusive agent, offering a fee of h^* , each buyer-exclusive agent accepts every invitation with an offer greater than or equal to h^* , and, play continues in the sell-side \hat{a} -collusive punishment phase (instead of the two-sided \hat{a} -collusive punishment phase) as appropriate.

The Bertrand reversion phase proceeds as expected, with each buyer-exclusive agent announcing a price of 0 on the buy side, each seller-proficient agent announcing a price of 0

on the sell side, buyers and sellers allocating themselves optimally given subsequent network formation (given prices), and statically optimal network formation.

The proof proceeds as in Theorem 1 *mutatis mutandis* except for the differences described in Section 4.2.2.

C.5 Proof of Theorem 4

As in our prior proofs, to simplify the exposition, we first define distinguished actions for buyers and sellers:

1. A buyer *lists seller-proficient agents arbitrarily* by reporting each ranking with equal probability that includes every seller-proficient agent; similarly, a seller *lists seller-proficient agents arbitrarily* by reporting each ranking with equal probability that includes every seller-proficient agent.
2. A buyer *prioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a first; similarly, a seller *prioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a first.
3. A buyer *deprioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a last; similarly, a seller *deprioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a last.

We further define distinguished actions for agents in the network formation steps:

1. The *full network among seller-proficient agents forms* when each seller-proficient agent a invites every other seller-proficient agent \tilde{a} , and every seller-proficient agent \tilde{a} accepts every invitation.
2. The *network among seller-proficient agents excluding a forms* when:
 - Each seller-proficient agent other than a does not invite a and invites every other seller-proficient agent \tilde{a} .
 - Agent a invites every other agent.
 - Each seller-proficient agent \tilde{a} other than a accepts every invitation he receives except invitations from a .

- Agent a accepts an invitation from any other agent.

Thus, the full network forms among seller-proficient agents other than a and no (seller-proficient) agent forms any links with a .

3. *No network forms* when no (seller-proficient) agent invites any other (seller-proficient) and a (seller-proficient) agent a accepts an invitation if and only if $p_B^a - c \geq 0$.

We also define the buyer deviation price in the cooperation phase as $p_B^\circ = v_B - \frac{1}{\kappa_S}(v_B - p_B^*)$ and the seller deviation price $p_S^\circ = v_S - \frac{1}{\kappa_B}(v_S - p_S^*) = v_S$; these are the prices at which buyers and sellers will be willing to work with a deviating agent in the cooperation phase. During the collusive punishment phase, prices are $q_B^* = (1 - \kappa_S)v_B - \kappa_S v_S$ and $q_S^* = v_S$. We define the buyer deviation price during the collusive punishment phase as $q_B^\circ = v_B - \frac{1}{\kappa_S}(v_B - q_B^*)$ and the seller deviation price as $q_S^\circ = v_S - \frac{1}{\kappa_B}(v_S - q_S^*) = v_S$; these are the prices at which buyers and sellers will be willing to work with an agent who deviates in the collusive punishment phase. Finally, we suppress here detailing the strategies after mutual deviations, i.e., after two or more agents *simultaneously* deviate: since no agent expects any other agent to deviate, such cases have no effect on incentives.

As in Theorem 1, the strategy profile that sustains p_B^* and p_S^* consists of three phases: In the *cooperation phase*:

1. Every seller-proficient agent offers a buyer price p_B^* and a seller price p_S^* and buyer-exclusive agents do not make offers.
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

Case 1: *Collusive pricing*: Each seller-proficient agent has offered (p_B^*, p_S^*) . Buyers and sellers list seller-proficient agents arbitrarily.

Case 2: *Ineffective price deviation by $\hat{a} \in A_S$* : Each seller-proficient agent except \hat{a} has offered (p_B^*, p_S^*) and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \not\leq (p_B^\circ, p_S^\circ)$. Buyers and sellers deprioritize agent \hat{a} among seller-proficient agents.

Case 3: *Effective price deviation by $\hat{a} \in A_S$* : Each agent except \hat{a} has offered (p_B^*, p_S^*) and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (p_B^\circ, p_S^\circ)$. Buyers and sellers prioritize agent \hat{a} among seller-proficient agents.

Case 4: *Price deviation by $\check{a} \in A_B$* : Each seller-proficient agent has offered (p_B^*, p_S^*) and a buyer-exclusive agent \check{a} has offered a price of $p_B^{\check{a}}$. Buyer list seller-proficient agents arbitrarily (and, in particular, do not rank \check{a}) and sellers list seller-proficient agents arbitrarily.

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing. The full network among seller-proficient agents forms.

Cases 2 and 3: Price deviation by $\hat{a} \in A_S$. If $p_B^* - c \geq 0$ then the network among seller-proficient agents excluding a forms; otherwise, no network forms.

Case 4: Price deviation by $\check{a} \in A_B$. The full network among seller-proficient agents forms.

4. Under collusive pricing or a price deviation by a buyer-exclusive agent, if the network implied by equilibrium play forms, play continues in the cooperation phase. After a price deviation by a seller-proficient agent \hat{a} , if the network implied by equilibrium play forms, play proceeds to the \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the \hat{a} -collusive punishment phase:

1. Every seller-proficient agent other than \hat{a} offers a buyer price q_B^* and a seller price q_S^* and buyer-exclusive agents do not make offers.
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing: Each seller-proficient agent has offered (p_B^*, p_S^*) . Buyers and sellers deprioritize agent \hat{a} among seller-proficient agents.

Case 2: Ineffective price deviation by $\hat{a} \in A_S$: Each seller-proficient agent except \hat{a} has offered (q_B^*, q_S^*) and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \not\leq (q_B^{\circ}, q_S^{\circ})$. Buyers and sellers deprioritize agent \hat{a} among seller-proficient agents.

Case 3: Effective price deviation by $\hat{a} \in A_S$: Each agent except \hat{a} has offered (q_B^*, q_S^*) and \hat{a} has offered $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (q_B^{\circ}, q_S^{\circ})$. Buyers and sellers prioritize agent \hat{a} among seller-proficient agents.

Case 4: Price deviation by $\check{a} \in A_B$: Each seller-proficient agent has offered (q_B^*, q_S^*) and a buyer-exclusive agent \check{a} has offered a price of $p_B^{\check{a}}$. Buyer list seller-proficient agents arbitrarily (and, in particular, do not rank \check{a}) and sellers list seller-proficient agents arbitrarily.

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing. The network among seller-proficient agents excluding a forms.

Cases 2 and 3: Price deviation by $\hat{a} \in A_S$. If $q_B^* - c \geq 0$ then the network among seller-proficient agents excluding \hat{a} forms; otherwise, no network forms.

Case 4: Price deviation by $\check{a} \in A_B$. The full network among seller-proficient agents forms.

4. Under collusive pricing or a price deviation by a buyer-exclusive agent, if the network implied by equilibrium play forms, play continues in the \hat{a} -collusive punishment phase. After a price deviation by a seller-proficient agent \hat{a} , if the network implied by equilibrium play forms, play proceeds to the \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the Bertrand reversion phase, each buyer-exclusive agent offers a price of 0 to buyers and each seller-proficient agent offers a price of 0 to sellers. Buyers rank all buyer-exclusive agents offering non-negative prices, with those agents offering lower prices ranked higher; similarly, sellers rank all seller-proficient agents offering non-negative prices, with those agents offering lower prices ranked higher. Finally, the full network forms.

The proof then proceeds as in Theorem 3 *mutatis mutandis* except for the differences described in Section 4.2.3.

C.6 Proof of Theorem 5

We will show the highest sustainable profits when buyers are represented by buyer-exclusive agents are found by solving

$$\max_{p_B^*, p_S^*} \{p_B^* + p_S^*\} \tag{5}$$

subject to the constraint that no seller-proficient agent can profitably deviate by attracting both buyers and sellers and not working other agents,

$$\frac{\sigma}{1 - \delta} (p_B^* + p_S^*) \geq \left(\left(v_B - \frac{1}{\kappa_S} (v_B - p_B^*) \right) + \left(v_S - \frac{1}{\kappa_B} (v_S - p_S^*) \right) \right) \kappa_B \kappa_S,$$

the individual rationality constraints for the buyers and sellers that $p_B^* \leq v_B$ and $p_S^* \leq v_S$, and the individual rationality constraint for buyer-exclusive agents that $p_B^* \geq 0$. Solving this linear program yields the price of Theorem 5.

As in our prior proofs, to simplify the exposition, we first define distinguished actions for buyers and sellers:

1. A buyer *lists buyer-exclusive agents arbitrarily* by reporting each ranking with equal probability that includes every buyer-exclusive agent; similarly, a seller *lists seller-proficient agents arbitrarily* by reporting each ranking with equal probability that includes every seller-proficient agent.
2. A buyer *prioritizes agent a among buyer-exclusive agents* by reporting with equal probability each ranking that both includes every buyer-exclusive agent and ranks a first; similarly, *prioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a first.
3. A buyer *deprioritizes agent a among buyer-exclusive agents* by reporting with equal probability each ranking that both includes every buyer-exclusive agent and ranks a last; similarly, *deprioritizes agent a among seller-proficient agents* by reporting with equal probability each ranking that both includes every seller-proficient agent and ranks a last.

We further define distinguished actions for agents in the network formation steps:

1. The *regular network forms* when each seller-proficient agent a invites every buyer-exclusive agent \tilde{a} , and every buyer-exclusive agent \tilde{a} accepts every invitation.
2. The *network excluding (the seller-proficient agent) a forms* when:
 - Each seller-proficient agent other than a does not invite a and invites every buyer-exclusive agent \tilde{a} .
 - Agent a invites every other agent.
 - Each buyer-exclusive agent \tilde{a} accepts every invitation he receives except invitations from a .
 - Agent a accepts an invitation from any other agent.

Thus, the regular network forms among buyer-exclusive agents and seller-proficient agents other than a .

3. The *network excluding (the buyer-exclusive agent) a forms* when:
 - Each seller-proficient agent invites every buyer-exclusive agent except a .
 - Each buyer-exclusive agent accepts every invitation he receives.

Thus, the regular network forms among buyer-exclusive agents other than a and seller-proficient agents.

We also define the buyer deviation price in the cooperation phase as $p_B^\circ = v_B - \frac{1}{\kappa_S}(v_B - p_B^*)$ and the seller deviation price $p_S^\circ = v_S - \frac{1}{\kappa_B}(v_S - p_S^*) = v_S$; these are the prices at which buyers and sellers will be willing to work with a deviating seller-proficient agent in the cooperation phase. During the collusive punishment phase, prices are depend on whether $\hat{\sigma} \leq \kappa_S(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1-\kappa_S}{\kappa_S})$: We have that $q_B^* = \kappa_S v_B$ and $q_S^* = \kappa_B v_S$ when $\hat{\sigma} \geq \kappa_S(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1-\kappa_S}{\kappa_S})$; we have that $q_B^* = 0$ and $q_S^* = v_B \frac{\kappa_B}{\kappa_S}(1-\kappa_S) + v_S(1-\kappa_B) + \kappa_B c$ when $\hat{\sigma} \leq \kappa_S(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1-\kappa_S}{\kappa_S})$. We define the buyer deviation price during the collusive punishment phase as $q_B^\circ = v_B - \frac{1}{\kappa_S}(v_B - q_B^*)$ and the seller deviation price as $q_S^\circ = v_S - \frac{1}{\kappa_B}(v_S - q_S^*) = v_S$; these are the prices at which buyers and sellers will be willing to work with an agent who deviates in the collusive punishment phase. Finally, we suppress here detailing the strategies after mutual deviations, i.e., after two or more agents *simultaneously* deviate: since no agent expects any other agent to deviate, such cases have no effect on incentives.

As elsewhere, the strategy profile that sustains p_B^* and p_S^* consists of three phases: In the *cooperation phase*:

1. Every buyer-exclusive agent offers a buyer price p_B^* and every seller-proficient agent offers a seller price p_S^* .
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing: Each buyer-exclusive agent has offered a buyer price p_B^* and each seller-proficient agent has offered a seller price p_S^* . Buyers list buyer-exclusive agents arbitrarily and sellers list seller-proficient agents arbitrarily.

Case 2: Price deviation by a buyer-exclusive agent $\check{a} \in A_B$: Each buyer-exclusive agent except \check{a} has offered a buyer price p_B^* , the buyer-exclusive agent \check{a} has offered a price $p_B^{\check{a}} \neq p_B^*$, and each seller-proficient agent has offered a seller price p_S^* . Buyers deprioritize agent \check{a} among buyer-exclusive agents and sellers list seller-proficient agents arbitrarily.

Case 3: Ineffective price deviation by $\hat{a} \in A_S$: Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent other than \hat{a} has offered a seller price p_S^* , and agent \hat{a} has offered prices $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \not\leq (p_B^\circ, p_S^\circ)$. Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.

Case 4: Effective price deviation by $\hat{a} \in A_S$: Each buyer-exclusive agent has offered a buyer price p_B^* , each seller-proficient agent other than \hat{a} has offered a

seller price p_S^* , and agent \hat{a} has offered prices $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (p_B^*, p_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (p_B^{\circ}, p_S^{\circ})$. Buyers prioritize \hat{a} among buyer-exclusive agents and sellers prioritize \hat{a} among seller-proficient agents.

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing. The regular network forms.

Case 2: Price deviation by a buyer-exclusive agent $\check{a} \in A_B$. The network excluding the buyer-exclusive agent \check{a} forms.

Cases 3 and 4: Price deviation by $\hat{a} \in A_S$. The network excluding the seller-proficient agent \hat{a} forms.

4. Under collusive pricing or a price deviation by a buyer-exclusive agent, if the network implied by equilibrium play forms, play continues in the cooperation phase. After a price deviation by a seller-proficient agent \hat{a} , if the network implied by equilibrium play forms, play proceeds to the \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the \hat{a} -collusive punishment phase:

1. Every buyer-exclusive agent offers a buyer price q_B^* and every seller-proficient agent offers a seller price q_S^* .
2. Buyer and seller behavior depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing: Each buyer-exclusive agent has offered a buyer price q_B^* and each seller-proficient agent has offered a seller price q_S^* . Buyers list buyer-exclusive agents arbitrarily and sellers list seller-proficient agents arbitrarily.

Case 2: Price deviation by a buyer-exclusive agent $\check{a} \in A_B$: Each buyer-exclusive agent except \check{a} has offered a buyer price q_B^* , the buyer-exclusive agent \check{a} has offered a price $p_B^{\check{a}} \neq q_B^*$, and each seller-proficient agent has offered a seller price q_S^* . Buyers deprioritize agent \check{a} among buyer-exclusive agents and sellers list seller-proficient agents arbitrarily.

Case 3: Ineffective price deviation by $\hat{a} \in A_S$: Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent other than \hat{a} has offered a seller price q_S^* , and agent \hat{a} has offered prices $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \not\leq (q_B^{\circ}, q_S^{\circ})$. Buyers list buyer-exclusive agents arbitrarily and sellers deprioritize \hat{a} among seller-proficient agents.

Case 4: Effective price deviation by $\hat{a} \in A_S$: Each buyer-exclusive agent has offered a buyer price q_B^* , each seller-proficient agent other than \hat{a} has offered a seller price q_S^* , and agent \hat{a} has offered prices $(p_B^{\hat{a}}, p_S^{\hat{a}}) \neq (q_B^*, q_S^*)$ such that $(p_B^{\hat{a}}, p_S^{\hat{a}}) \leq (q_B^*, q_S^*)$. Buyers prioritize \hat{a} among buyer-exclusive agents and sellers prioritize \hat{a} among seller-proficient agents.

3. The invitations made and accepted also depend on the pricing behavior of the agents in Step 1:

Case 1: Collusive pricing. The regular network forms.

Case 2: Price deviation by a buyer-exclusive agent $\check{a} \in A_B$. The network excluding the buyer-exclusive agent \check{a} forms.

Cases 3 and 4: Price deviation by $\hat{a} \in A_S$. The network excluding the seller-proficient agent \hat{a} forms.

4. Under collusive pricing or a price deviation by a buyer-exclusive agent, if the network implied by equilibrium play forms, play continues in the \hat{a} -collusive punishment phase. After a price deviation by a seller-proficient agent \hat{a} , if the network implied by equilibrium play forms (i.e., \hat{a} is the only player excluded from the network), play proceeds to the \hat{a} -collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the Bertrand reversion phase, each buyer-exclusive agent offers a price of 0 to buyers and each seller-proficient agent offers a price of 0 to sellers. Buyers rank all buyer-exclusive agents offering non-negative prices, with those agents offering lower prices ranked higher; similarly, sellers rank all seller-proficient agents offering non-negative prices, with those agents offering lower prices ranked higher. Finally, the regular network forms.

The solution of the program (5) depends on whether $\hat{\sigma} \leq \kappa_S(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1-\kappa_S}{\kappa_S})$. There are two cases:

1. If $\hat{\sigma} \leq \kappa_S(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1-\kappa_S}{\kappa_S})$, it is immediate that buyer-exclusive agents are willing to exclude a price-deviating seller since they receive 0 payoffs regardless of their actions. Moreover, to attract buyers, a price-deviator would have to offer a buyer price less than 0, and so would not be willing to take his buyers to other agents' sellers in any case. Thus, a price-deviator will not work with any other agents.

The proof then proceeds as in Theorem 4 *mutatis mutandis*.

2. If $\hat{\sigma} > \kappa_S(1 - \frac{c}{v_S} - \frac{v_B}{v_S} \frac{1-\kappa_S}{\kappa_S})$, then recall that $q_B^* = \kappa_S v_B$ and $q_S^* = \kappa_B v_S$, and we can calculate $p_B^{\circ} = v_B - \frac{1}{\kappa_S}(v_B - q_B^*)$ and $p_S^{\circ} = v_S - \frac{1}{\kappa_B}(v_S - q_S^*)$.

In this case, both prices are positive, and so barring dynamic incentives agents will be tempted to work with a price deviator. Thus, we need to check that seller-proficient agents will not work with a buyer-exclusive agent if he deviates on price. In the cooperation phase, the payoff to a seller-proficient agent from refusing to invite a price-deviating buyer-exclusive agent is

$$\frac{\delta}{1-\delta}\sigma p_S^*$$

while the payoff from inviting the price-deviating buyer-exclusive agent is

$$\kappa_B \sigma p_S^*$$

and so it is sufficient that $\delta \geq \frac{1}{2}$. A similar calculation holds in the collusive punishment phase.

We also need to check that if a seller-proficient agent deviates on prices both buyer-exclusive and seller-proficient agents will not work with the deviator. To see that buyer-exclusive agents will not work with a deviator, note that a buyer-exclusive agent gets

$$\frac{\delta}{1-\delta}\beta q_B^* = \frac{\delta}{1-\delta}\beta \kappa_S v_B$$

by refusing to work with the deviator, but at most

$$\beta \kappa_S p_B^* \leq \beta \kappa_S v_B$$

by working with the deviator; thus for $\delta \geq \frac{1}{2}$ it is optimal to refuse an invitation from the deviator. To see that seller-proficient agents will not work with a deviator, note that a seller-proficient agent gets

$$\frac{\delta}{1-\delta}\sigma q_S^* = \frac{\delta}{1-\delta}\sigma \kappa_B v_S$$

by refusing to work with the deviator, but at most

$$\sigma \kappa_B p_S^* = \sigma \kappa_B v_S$$

by working with the deviator; thus for $\delta \geq \frac{1}{2}$ it is optimal to refuse an invitation from the deviator.

Finally, we need to show that at prices (q_B^*, q_S^*) no agent will wish to offer prices of (q_B^0, q_S^0) —in particular, that an agent getting 0 in the collusive punishment phase will

not wish to offer prices of (q_B°, q_S°) . To see this, we calculate that

$$\begin{aligned}
\kappa_B \kappa_S (q_B^\circ + q_S^\circ - c) &= \kappa_B \kappa_S \left(v_B - \frac{1}{\kappa_S} (v_B - q_B^*) + v_S - \frac{1}{\kappa_B} (v_S - q_S^*) - c \right) \\
&= \kappa_B \kappa_S \left(v_B - \frac{1}{\kappa_S} (v_B - \kappa_S v_B) + v_S - \frac{1}{\kappa_B} (v_S - \kappa_B v_S) - c \right) \\
&= \kappa_B v_B (2\kappa_S - 1) + \kappa_S v_S (2\kappa_B - 1) - \kappa_B \kappa_S c \\
&\leq 0.
\end{aligned}$$

A similar calculation holds in the collusive punishment phase.

The proof then proceeds as in Theorem 4 *mutatis mutandis*.

D Analyzing Asymmetric Equilibria in the Stage Game

In the main construction of our equilibrium, we have assumed that buyers and sellers treat identically agents who are treated identically by other agents (with respect to network formation). Suppose, by contrast, that buyers and sellers allocated themselves selectively following a price deviation not only by working preferentially with the price deviator but also with another non-deviating agent; the buyers and sellers could then potentially unravel collusion by incentivizing said non-deviating agent to work with the price deviator. For example, if all buyers and sellers coordinated on listing the price deviator first and a particular non-deviating agent a second following a price deviation, then a would fill to capacity; thus a would have a much stronger temptation to invite and/or accept an invitation from the price deviator. Indeed, given the continuation play that we use in the construction of our equilibrium, such an agent a would deviate from the equilibrium to work with the price deviator.

In fact, buyer-and-seller coordination-proofness would imply that this behavior would occur. A set of buyers and sellers could coordinate to work with a price deviator and a subset of non-price deviators, incentivizing those non-price deviators to work with the price deviator. However, a minor modification to the analysis in the text rules this out; regardless of how buyers and sellers allocate themselves among non-deviating agents, future payoffs can be constructed such that all non-deviating agents exclude the price deviator. To do this, we must condition future payoffs to each non-deviating agent on his temptation to work with the price deviator—we do this by allocating sellers in future periods to each agent based on that agent’s temptation to work the price deviator today.

First, consider a price deviation in the cooperation phase. We show that an agent a is better off following his prescribed actions than if he invited \hat{a} with a fee of $-p_B^{\hat{a}}$ and accepted

an invitation from \hat{a} with a fee of $p_S^{\hat{a}}$ (and followed his prescribed actions with respect to other agents). Let $\tilde{\kappa}_B^{\circ}$ be the share of buyers that selected the deviating agent \hat{a} and let $\tilde{\kappa}_S^{\circ}$ be the share of sellers that selected \hat{a} . Let $\tilde{\kappa}_B^a$ be the share of buyers that selected a non-deviating agent a and let $\tilde{\kappa}_S^a$ be the share of sellers that selected a . Finally, let κ_B^a be the share of sellers allocated to agent a in future collusive play. The total payoff for a from following his prescribed actions is

$$\underbrace{\underbrace{\tilde{\kappa}_S^a}_{\text{Mass of sellers represented by } a} \underbrace{(1 - \tilde{\kappa}_B^{\circ})}_{\text{Mass of buyers represented by agents other than } \hat{a}} \underbrace{(p_B^* + p_S^*)}_{\text{Profit per transaction}}}_{\text{Profits from working with agents other than } \hat{a} \text{ this period}} + \underbrace{\frac{\delta}{1 - \delta} \kappa_S^a (q_B^* + q_S^*)}_{\text{Payoff in future periods from adhering}} \quad (6)$$

while the payoff from working with \hat{a} is at most

$$\underbrace{\tilde{\kappa}_S^a (1 - \tilde{\kappa}_B^{\circ}) (p_B^* + p_S^*)}_{\text{Profits from working with agents other than } \hat{a} \text{ this period}} + \underbrace{\underbrace{\tilde{\kappa}_S^a}_{\text{Mass of sellers represented by } a} \underbrace{\tilde{\kappa}_B^{\circ}}_{\text{Mass of buyers represented by } \hat{a}} \underbrace{(p_B^{\hat{a}} + p_S^*)}_{\text{Profit per transaction}} + \underbrace{\tilde{\kappa}_B^a}_{\text{Mass of buyers represented by } a} \underbrace{\tilde{\kappa}_S^{\circ}}_{\text{Mass of sellers represented by } \hat{a}} \underbrace{(p_B^* + p_S^{\hat{a}})}_{\text{Profit per transaction}}}_{\text{Payoff from working with } \hat{a}} \quad (7)$$

Since the payoff to working with non-deviating agents is identical regardless of the decision to work with the deviating agent, we have the condition for adhering as

$$\frac{\delta}{1 - \delta} \kappa_S^a (q_B^* + q_S^*) \geq \tilde{\kappa}_S^a \tilde{\kappa}_B^{\circ} (p_B^{\hat{a}} + p_S^*) + \tilde{\kappa}_B^a \tilde{\kappa}_S^{\circ} (p_B^* + p_S^{\hat{a}})$$

Observe the following:

- Since $q_B^* = (1 - \kappa_S)v_B - \kappa_S v_S$ and $q_S^* = v_S$, we have that $q_B^* + q_S^* = (v_B + v_S)(1 - \kappa_S)$.
- Since $p_B^{\hat{a}} \leq p_B^{\circ} \leq p_B^* \leq v_B$ and $p_S^* = v_S$, we have that $p_B^{\hat{a}} + p_S^* \leq v_B + v_S$.
- Since $p_S^{\hat{a}} \leq p_S^{\circ} = p_S^* = v_S$ and $p_B^* \leq v_B$, we have that $p_B^* + p_S^{\hat{a}} \leq v_B + v_S$.

Hence, it is sufficient that

$$\frac{\delta}{1 - \delta} \kappa_S^a (1 - \kappa_S) \geq \tilde{\kappa}_B^{\circ} \tilde{\kappa}_S^a + \tilde{\kappa}_B^a \tilde{\kappa}_S^{\circ} \quad (8)$$

for each agent a other than \hat{a} .

Thus, to ensure that each agent is incentivized to not work with the price deviator, we need to find κ_S^a for each agent a so that each $\kappa_S^a \leq \kappa_S$ and that $\sum_{a \in A \setminus \{\hat{a}\}} \kappa_S^a = 1$.

To see that $\kappa_S^a \leq \kappa_S$, note that the right-hand side (8) is maximized when $\tilde{\kappa}_B^\circ = \tilde{\kappa}_B^a = \kappa_B \leq \frac{1}{3}$ and $\tilde{\kappa}_S^\circ = \tilde{\kappa}_S^a = \kappa_S \leq \frac{1}{3}$. Thus, it is enough that

$$\frac{\delta}{1-\delta} \kappa_S^a (1 - \kappa_S) \geq \frac{2}{9};$$

this can always be satisfied for $\delta \geq \frac{1}{2}$ by setting $\kappa_S^a = \frac{1}{3}$.

To see that $\sum_{a \in A \setminus \{\hat{a}\}} \kappa_S^a \leq 1$, we sum (8) over all $a \in A \setminus \{\hat{a}\}$ to obtain the requirement that

$$\begin{aligned} \frac{\delta}{1-\delta} \sum_{a \in A \setminus \{\hat{a}\}} \kappa_S^a &= \frac{\tilde{\kappa}_B^\circ \left(\sum_{a \in A \setminus \{\hat{a}\}} \tilde{\kappa}_S^a \right) + \tilde{\kappa}_S^\circ \left(\sum_{a \in A \setminus \{\hat{a}\}} \tilde{\kappa}_B^a \right)}{1 - \kappa_S} \\ &\geq \frac{\tilde{\kappa}_B^\circ (1 - \tilde{\kappa}_S^\circ) + \tilde{\kappa}_S^\circ (1 - \tilde{\kappa}_B^\circ)}{1 - \kappa_S} \end{aligned} \quad (9)$$

as only buyers and sellers who do not work with the price deviator can work with the other agents. Note that (9) is maximized when $\tilde{\kappa}_B^\circ = \kappa_B$ and $\tilde{\kappa}_S^\circ = \kappa_S$ as we must have that $\tilde{\kappa}_B^\circ < \frac{1}{2}$ and $\tilde{\kappa}_S^\circ < \frac{1}{2}$. Thus, we need that

$$\frac{\delta}{1-\delta} \sum_{a \in A \setminus \{\hat{a}\}} \kappa_S^a \geq \frac{\kappa_B (1 - \kappa_S) + \kappa_S (1 - \kappa_B)}{1 - \kappa_S}.$$

Thus, since $\kappa_B \geq \kappa_S$, it is sufficient that

$$\frac{\delta}{1-\delta} \sum_{a \in A \setminus \{\hat{a}\}} \kappa_S^a \geq \kappa_B + \kappa_S.$$

Hence, we can find κ_S^a that incentivize each non-deviating agent to not work with the price deviator so long as $\delta \geq \frac{1}{2}$ as $\kappa_S \leq \kappa_B \leq \frac{1}{3}$.

An analogous construction can incentivize agents to not work with a price deviator in the collusive punishment phase.