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

January 14, 2019

PRELIMINARY AND INCOMPLETE:
PLEASE DO NOT CITE OR DISTRIBUTE WITHOUT PERMISSION

Abstract

The U.S. residential real estate agency market presents a puzzle for economic theory: agent entry is frequent and agents' costs of providing service are low, yet commissions on real estate transactions have remained constant and high for decades. We model the real estate agency market, and other brokered markets, as a repeated extensive form game; in our game, brokers first post prices for customers and then choose which agents on the other side of the market to work with. We show that prices appreciably higher than the competitive prices can be sustained (for a fixed discount factor) regardless of the number of brokers through strategies that condition willingness to transact with each broker on that broker's initial posted prices. Our results can thus rationalize why this market exhibits both fierce competition for customers and pricing high above marginal cost; moreover, our model can help explain why agents and platforms who have tried to reduce commissions have had trouble entering the market.

JEL Classification: D43, L13, L4, R39

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

*The authors thank Panle Jia Barwick, Michael Ostrovsky, Maisy Wong, and several seminar audiences for helpful comments. 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 (Delcoure and Miller, 2002). Yet the market for providing residential real estate brokerage services is quite easy to enter, and so we might naturally expect price competition to drive down fees; 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".[1]

We provide a potential explanation for how the real estate brokerage industry maintains high prices even in the presence of many independent brokers. 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 price for intermediation services; each buyer and seller then chooses an agent 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).

[1]Similarly, Levitt and Syverson (2008) 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).

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.

Of course, it must be incentive compatible for agents to refuse to work with a price deviator. Here, we rely on the repeated interactions between agents: A non-price deviating agent is willing to forego working with a price deviator today if he is sufficiently rewarded for doing so in the future. First, the non-price deviating agent's foregone 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 no-profit equilibrium; 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 forego working with a price deviator today for reasonably high discount factors.

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. Our setting is a repeated extensive form game and, as noted by Mailath et al. (2017), such games may require more involved responses to deviations because the need to reward within-period punishments implies that rewarding agents in future periods may be more important than continuing to punish the initial deviator.[2] This is the case in our model: we cannot simply revert to the no-profit equilibrium after an agent undercuts on prices as we need to reward other agents for forgoing working with that price deviator.

Our analysis not only explains how real estate agents may maintain high prices, but also

[2]It 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.

enables us to assess different policy responses that have been suggested. Han and Hong (2011) investigated "rebate bans," laws that prohibit buyers' agents from sharing their commission with buyers (although agents are still allowed to pay buyers' closing costs); ten states have such a policy. We show that rebate bans facilitate higher agency fees 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 finding even accords with surprisingly candid remarks of the real estate agents themselves.[3]

Meanwhile, Barwick (2018) has suggested banning "agency fees," i.e., the commission that a seller agent pays to a buyer agent upon the completion of a transaction; as Barwick (2018) notes, countries that have adopted this policy have lower overall agency fees (even though fees for buyers increase). In our framework, the optimal collusive scheme involves fully exploiting sellers while possibly charging buyers less than cost. Eliminating agency fees make charging buyers less than cost non-viable, 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 falls.

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, and Section 4.2 considers the effects of eliminating agency fees on prices. Section 5 concludes.

[3]The 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.

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 *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 κ_B and κ_S are both 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 discount the future at rate $\delta \in (0,1)$.

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* $p^a_{B,t} \in \mathbb{R}$ and a *seller price* $p^a_{S,t} \in \mathbb{R}$. All prices are publicly observed.

Step 2: Each buyer $b \in B_t$ selects a subset of agents and submits an ordered list over that subset; that is, each buyer reports preferences 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.[4] 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 $s \in S_t$ selects a subset of agents and submits an ordered list over that subset; that is, each

[4]This 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 preference lists in Appendix B.

seller reports preferences 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.

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, each accepted invitation along with its associated fee is publicly observed.

We can think of the set of accepted invitations as generating a *(directed) network*, where a link from $\tilde{a}$ to a represents the fact that $\tilde{a}$ has accepted a's invitation. We say that a *complete network forms* among agents $\bar{A} \subseteq A$ 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 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 a match 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 $\left(v_B - p_{B,t}^{\mathbf{a}(b)}\right)|\cup_{a\in A^{\mathbf{a}(b)\Rightarrow}} \mathbf{S}_t(a)|$. That is, the payoff is the net value of a purchase times the measure of connected sellers, which gives the probability of a transaction.

2. The expected payoff for seller s in period t is given by $\left(v_S - p_{S,t}^{\mathbf{a}(s)}\right)|\cup_{a\in A^{\Rightarrow \mathbf{a}(s)}} \mathbf{B}_t(a)|$. That is, the payoff is the net value of a sales times the measure of connected buyers, which gives the probability of a transaction.

3. The expected payoff for an agent a in period t is given by

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

   this reduces to

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

   the total revenue from buy and sell side transactions.

2.2 Equilibrium Definition

In our setting, perfect collusion can 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. This is achieved by having both buyers and sellers refuse to sign up with any agent if prices are not as expected; this 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; this is a version of the classic indeterminacy of equilibrium in platforms result (Weyl, 2010).

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 pathological outcomes, we thus 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, 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 buyer and sellers as a mechanism to support higher prices.

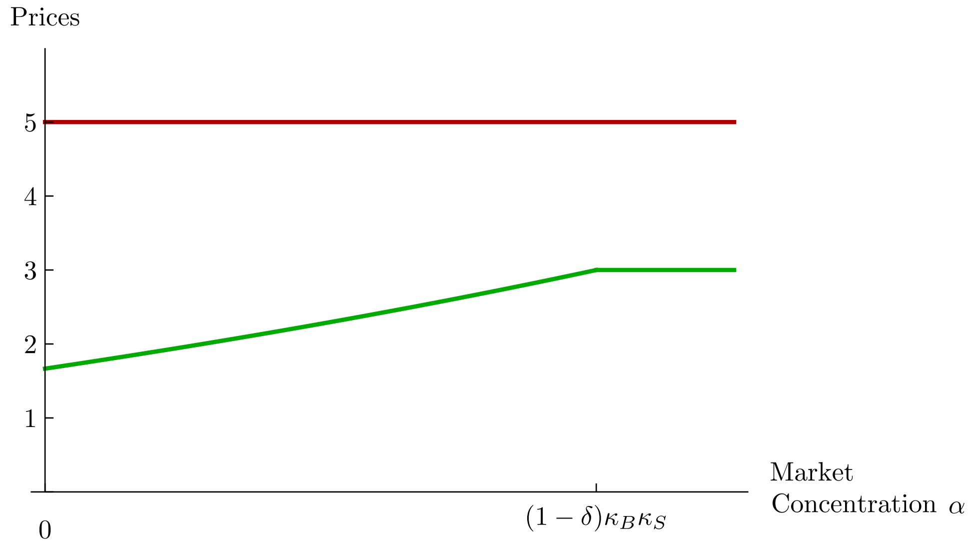

Figure 1: The prices supporting the highest sustainable profit equilibrium. The dark red line is the price agents charge to sellers, $p_S^\star$, and the light green line is the price agents charge to buyers, $p_B^\star$. Here, $\delta = \frac{3}{4}$, $v_B = 3$, $v_S = 5$, $\kappa_B = \frac{1}{5}$, and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

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*

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

$$p_S^\star = v_S.$$

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

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 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 highest sustainable profits equilibrium would require setting the buyer price to v_B and the seller price below v_S. We now move on to constructing the equilibrium, 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*.*

We now formally describe 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 preference ordering 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 preference ordering is reported with equal probability by each buyer.

   - Each seller reports a preference ordering 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 preference ordering 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).[5]

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

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

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). We call the actions with respect to making and accepting invitations described in this equilibrium the *statically optimal network formation actions*.

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 complete 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.[6] Thus, a can not increase his profits by increasing his buyer price or his seller price. It is immediate that a can not increase his profits by decreasing his buyer or seller price.

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 $(p_B^\star, p_S^\star)$ 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 $(p_B^\star, p_S^\star)$.

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 $p_B^\star$ and a seller price of $p_S^\star$. Assuming all the agents offer $p_B^\star$ and $p_S^\star$, buyers and sellers allocate themselves evenly across all agents; agents then

[6] Recall 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 $(|A| - 1)\kappa_S \geq 1$.

form the complete 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. Thus, if an agent undercuts on price on each side by a small amount, instead of increasing his market share—as one might expect—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 f 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 f. If all other agents exclude the price deviator f 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 $q_B^\star$ and a seller price of $q_S^\star$. Moreover, agents continue to exclude the price deviator f from the network. Note that, as in the cooperation phase, a (possibly new) price deviator must substantially undercut $(q_B^\star, q_S^\star)$ 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* if $(p_B, p_S) \leq (q_B^\circ, q_S^\circ)$. The prices $q_B^\star$ and $q_S^\star$ 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 for each agent other than the price deviator.

We now formally construct a strategy profile that sustains $(p_B^\star, p_S^\star)$.[7] Formally, 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:

   Case 1: ***Collusive pricing.*** In this case, each agent has offered $(p_B^\star, p_S^\star)$. Each buyer reports a preference ordering over all agents; each such preference ordering is reported with equal probability by each buyer. Similarly, each seller reports a preference ordering over all agents; each such preference ordering is reported with equal probability by each seller.

   Case 2: ***Ineffective price deviation by*** $\mathring{a}$**.** In this case, each agent except $\mathring{a}$ has offered $(p_B^\star, p_S^\star)$ and $\mathring{a}$ has offered $(p_B^{\mathring{a}}, p_S^{\mathring{a}}) \nleq (p_B^\circ, p_S^\circ)$, where the buyer deviation price p_B° is given by $v_B - \frac{1}{\kappa_S}(v_B - p_B^\star)$ and the seller deviation price p_S° is given by $v_S - \frac{1}{\kappa_B}(v_S - p_S^\star)$. Each buyer reports a preference ordering over all agents except $\mathring{a}$; each such preference ordering is reported with equal probability by each

[7] Here, we require that buyers and sellers treat agents with identical histories identically; that is, they do not discriminate between agents who have offered the same prices in every period (including the current one). This restriction prevents implausible coordination by buyers and sellers. In Appendix D, we relax this assumption and show that Theorem 1 still holds.

buyer. Each seller reports a preference ordering over all agents except $\mathring{a}$; each such preference ordering is reported with equal probability by each seller.

Case 3: ***Effective price deviation by*** $\mathring{a}$**.** In this case, each agent except $\mathring{a}$ has offered $(p_B^\star, p_S^\star)$ and $\mathring{a}$ has offered $(p_B^{\mathring{a}}, p_S^{\mathring{a}}) \leq (p_B^\circ, p_S^\circ)$. Each buyer reports a preference ordering over all agents with $\mathring{a}$ listed first; each such preference ordering is reported with equal probability by each buyer. Each seller reports a preference ordering over all agents with $\mathring{a}$ listed first; each such preference ordering is reported with equal probability by each seller.

Case 4: ***Mutual deviations.*** In this case, two or more agents have not offered $(p_B^\star, p_S^\star)$.[8] 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).

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

Case 1: Collusive pricing. Each agent invites every other agent, offering a fee of $-p_B^\star$, and every agent accepts every invitation with an offer greater than or equal to $-p_B^\star$.[9]

Cases 2 and 3: Price deviation by $\mathring{a}$**.** Each agent a other than $\mathring{a}$ does not invite $\mathring{a}$ and invites every other agent $\tilde{a}$ with a fee of $f^{\tilde{a}\leftarrow a} = -p_B^\star$. The agent $\mathring{a}$ invites every other agent a with a fee of $f^{a\leftarrow\mathring{a}} = p_S^{\mathring{a}}$. Each agent $\tilde{a}$ other than $\mathring{a}$ accepts every invitation he receives with a fee greater than or equal to $-p_B^\star$, except an invitation from $\mathring{a}$; the agent $\tilde{a}$ accepts an invitation from $\mathring{a}$ if and only if the fee is

[8]Note that this case requires simultaneous deviations by two or more agents and thus the payoffs in this case have no effect on any agent's incentives.

[9]The fee chosen here implies 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 strategy profile here to split the profits more evenly; in particular, any fee in $[-p_B^\star, p_S^\star]$ can be used instead to support $(p_B^\star, p_S^\star)$ but this requires verifying additional incentive constraints. Alternative fees in $[-p_B^\star, p_S^\star]$ between non-price deviators can also supported in Cases 2 and 3 of the cooperation phase, as well as Cases 1–5 of the collusive punishment phase.

greater than $p_S^{\mathring{a}}$. Agent $\mathring{a}$ accepts an invitation from $a \in A \setminus \{\mathring{a}\}$ if and only if the fee $f^{a \leftarrow \mathring{a}}$ is no less than $-p_B^{\mathring{a}}$.

Case 4: Mutual deviations. 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 all of the surplus that $\tilde{a}$ will receive from a buyer, unless that surplus is insufficient to compensate a for the negative price he is receiving from his seller. Each agent $\tilde{a}$ accepts the invitation from a as long as $f^{\tilde{a} \leftarrow a} \geq -p_B^{\tilde{a}}$.

4. Under collusive pricing, if the network (including fees) implied by equilibrium play forms, play continues in the cooperation phase. After a price deviation by $\mathring{a}$, if

   - the network (including fees) between agents other than $\mathring{a}$ implied by equilibrium play forms,
   - no agent invites $\mathring{a}$, and
   - each agent accepts an invitation from $\mathring{a}$ if and only if $f^{a \leftarrow \mathring{a}} > p_S^{\mathring{a}}$,

   then play proceeds to the $\mathring{a}$-collusive punishment phase. Otherwise, play proceeds to the Bertrand reversion phase.

In the *$\mathring{a}$-collusive punishment phase*:

1. Every agent (including $\mathring{a}$) offers buyer price $q_B^{\star} = (1 - \kappa_S)v_B - \kappa_S v_S$ and seller price $q_S^{\star} = v_S$.

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 $(q_B^{\star}, q_S^{\star})$. Each buyer reports a preference ordering over all agents except for $\mathring{a}$; each such preference ordering is reported with equal probability by each buyer. Similarly, each seller reports a preference ordering over all agents except for $\mathring{a}$; each such preference ordering is reported with equal probability by each seller.

Case 2: ***Ineffective price deviation by*** $\mathring{a}$**.** Each agent $a \in A \setminus \{\mathring{a}\}$ has offered $(q_B^\star, q_S^\star)$, and $\mathring{a}$ has offered $(p_B^{\mathring{a}}, p_S^{\mathring{a}})$ such that $(p_B^{\mathring{a}}, p_S^{\mathring{a}}) \neq (q_B^\star, q_S^\star)$ and $(p_B^{\mathring{a}}, p_S^{\mathring{a}}) \nleq (q_B^\circ, q_S^\circ)$, where the buyer deviation price q_B° is given by $v_B - \frac{1}{\kappa_S}(v_B - q_B^\star)$ and the seller deviation price q_S° is given by $v_S - \frac{1}{\kappa_B}(v_S - q_S^\star)$. Each buyer reports a preference ordering over all agents except for $\mathring{a}$; each such preference ordering is reported with equal probability by each buyer. Similarly, each seller reports a preference ordering over all agents except for $\mathring{a}$; each such preference ordering is reported with equal probability by each seller.

Case 3: ***Effective price deviation by*** $\mathring{a}$**.** Each agent $a \in A \setminus \{\mathring{a}\}$ has offered $(q_B^\star, q_S^\star)$, and $\mathring{a}$ has offered $(p_B^{\mathring{a}}, p_S^{\mathring{a}}) \leq (q_B^\circ, q_S^\circ)$. Each buyer reports a preference ordering over all agents with $\mathring{a}$ listed first; each such preference ordering is reported with equal probability by each buyer. Each seller reports a preference ordering over all agents with $\mathring{a}$ listed first; each such preference ordering is reported with equal probability by each seller.

Case 4: ***Ineffective price deviation by*** $\check{a} \neq \mathring{a}$**.** Each agent $a \in A \setminus \{\check{a}\}$ has offered $(q_B^\star, q_S^\star)$, and $\check{a}$ has offered $(p_B^{\check{a}}, p_S^{\check{a}}) \nleq (q_B^\circ, q_S^\circ)$. Each buyer reports a preference ordering over all agents except $\check{a}$ (including $\mathring{a}$); each such preference ordering is reported with equal probability by each buyer. Each seller reports a preference ordering over all agents except $\check{a}$ (including $\mathring{a}$); each such preference ordering is reported with equal probability by each seller.

Case 5: ***Effective price deviation by*** $\check{a} \neq \mathring{a}$**.** Each agent $a \in A \setminus \{\check{a}\}$ has offered $(q_B^\star, q_S^\star)$, and $\tilde{a}$ has offered $(p_B^{\check{a}}, p_S^{\check{a}}) \leq (q_B^\circ, q_S^\circ)$. Each buyer reports a preference ordering over all agents with $\check{a}$ listed first; each such preference ordering is reported with equal probability by each buyer. Each seller reports a preference ordering over all agents with $\check{a}$ listed first; each such preference ordering is reported with equal probability by each seller.

Case 6: Mutual deviations. Two or more agents have offered prices other than those prescribed by the equilibrium strategy profile.[10] 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).

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

Cases 1–3: Collusive pricing and price deviations by $\mathring{a}$.

Each agent a other than $\mathring{a}$ does not invite $\mathring{a}$ and invites every other agent $\tilde{a}$ with a fee of $f^{\tilde{a} \leftarrow a} = -p_B^{\star}$. The agent $\mathring{a}$ invites every other agent a with a fee of $f^{a \leftarrow \mathring{a}} = p_S^{\mathring{a}}$. Each agent $\tilde{a}$ other than $\mathring{a}$ accepts every invitation he receives with a fee greater than or equal to $-p_B^{\star}$, except an invitation from $\mathring{a}$; the agent $\tilde{a}$ accepts an invitation from $\mathring{a}$ if and only if the fee is greater than $p_S^{\mathring{a}}$. Agent $\mathring{a}$ accepts an invitation from $a \in A \smallsetminus \{\mathring{a}\}$ if and only if the fee $f^{a \leftarrow \mathring{a}}$ is no less than $-p_B^{\mathring{a}}$.

Cases 4 and 5: Effective and ineffective price deviations by $\check{a} \neq \mathring{a}$.

Each agent a other than $\check{a}$ does not invite $\check{a}$ and invites every other agent $\tilde{a}$ with a fee of $f^{\tilde{a} \leftarrow a} = -p_B^{\star}$. The agent $\check{a}$ invites every other agent a with a fee of $f^{a \leftarrow \check{a}} = p_S^{\check{a}}$. Each agent $\tilde{a}$ other than $\check{a}$ accepts every invitation he receives with a fee greater than or equal to $-p_B^{\star}$, except an invitation from $\check{a}$; the agent $\tilde{a}$ accepts an invitation from $\check{a}$ if and only if the fee is greater than $p_S^{\check{a}}$. Agent $\check{a}$ accepts an invitation from $a \in A \smallsetminus \{\check{a}\}$ if and only if the fee $f^{a \leftarrow \check{a}}$ is no less than $-p_B^{\check{a}}$.

Case 6: Mutual deviations. 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 all of the surplus that $\tilde{a}$ will receive from a buyer, unless that surplus is insufficient to compensate a for the negative price he is receiving from the seller. Each agent $\tilde{a}$ accepts the invitation from a as long as $f^{\tilde{a} \leftarrow a} \geq -p_B^{\tilde{a}}$.

[10] Note that this case requires simultaneous deviations by two or more agents and thus has no effect on any agent's expected payoffs from any action.

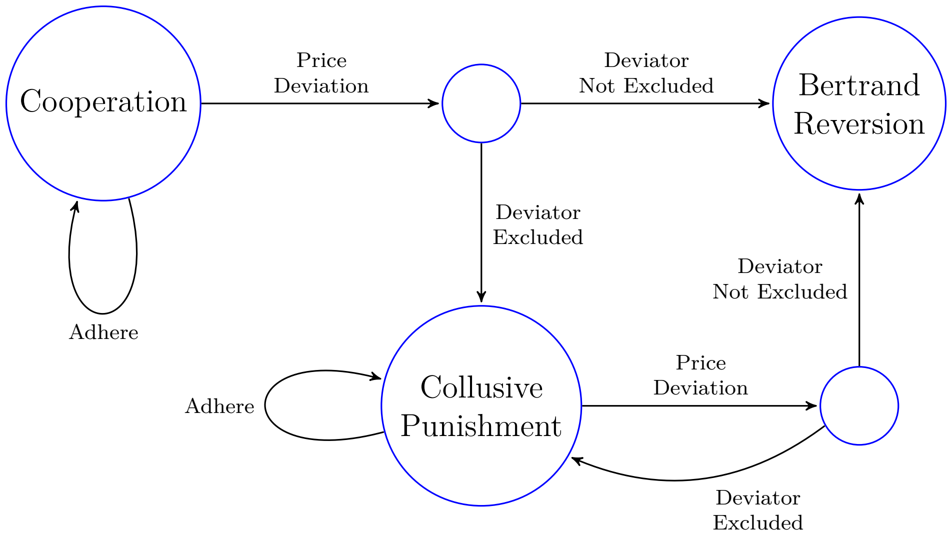

Figure 2: Automaton representation of the equilibrium we consider. Labeled nodes are phases; unlabeled nodes are intermediate phases, which represent the branching of transitions based on behavior in the later steps of the game.

4. In Cases 1–3, if

   - the network (including fees) between agents other than $\mathring{a}$ implied by equilibrium play forms,
   - no agent invites $\mathring{a}$, and
   - each agent accepts an invitation from $\mathring{a}$ if and only if $f^{a\leftarrow\mathring{a}} > p_S^{\mathring{a}}$,

   then play proceeds to the $\mathring{a}$-collusive punishment phase; otherwise, play proceeds to the Bertrand reversion phase. In Cases 4 and 5, if

   - the network (including fees) between agents other than $\check{a}$ implied by equilibrium play forms,
   - no agent invites $\check{a}$, and
   - each agent accepts an invitation from $\check{a}$ if and only if $f^{a\leftarrow\check{a}} > p_S^{\check{a}}$,

   then play proceeds to the $\check{a}$-collusive punishment phase; otherwise, play proceeds to the Bertrand reversion phase. In Case 6, 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 $(p_B^\star, p_S^\star)$ 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

There are four cases to consider:

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

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

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 $-p_B^\star$ (as this maximizes in-period profits). Otherwise, it is optimal for a to accept each offered fee of $-p_B^\star$ 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 $-p_B^\star$ as this maximizes in-period profits as well as ensuring that play continues in the cooperation phase.

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

Case 3: Effective price deviation by $\mathring{a}$. Showing that all agents other than $\mathring{a}$ follow prescribed play with respect to making invitations to and accepting invitations from agents other than $\mathring{a}$ is exactly as in the collusive pricing case. Moreover, it is immediate that $\mathring{a}$'s prescribed actions are optimal:

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

We now show the important fact that an agent a is better off following his prescribed actions than if he invited $\mathring{a}$ with a fee of $-p_B^{\mathring{a}}$ and accepted an invitation from $\mathring{a}$ with a fee of $p_S^{\mathring{a}}$ (and followed his prescribed actions with respect to other agents). It is then immediate that a is better off following his prescribed action than any other strategy by a with respect to inviting $\mathring{a}$ and/or accepting $\mathring{a}$'s invitation with a fee of $p_S^{\mathring{a}}$ (and actions with respect to other agents).

The total payoff for a from following his prescribed actions is[11]

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

while the payoff for working with $\mathring{a}$ is given by

$$\underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)(1-\kappa_B)(p_B^\star + p_S^\star)}_{\substack{\text{Profits from working with agents}\\ \text{other than } \mathring{a} \text{ this period}}} +$$
$$\underbrace{\underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)}_{\substack{\text{Mass of sellers}\\ \text{represented by } a}} \underbrace{\kappa_B}_{\substack{\text{Mass of buyers}\\ \text{represented by } \mathring{a}}} \underbrace{(p_B^{\mathring{a}} + p_S^\star)}_{\substack{\text{Profit per}\\ \text{transaction}}} + \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_B)}_{\substack{\text{Mass of buyers}\\ \text{represented by } a}} \underbrace{\kappa_S}_{\substack{\text{Mass of sellers}\\ \text{represented by } \mathring{a}}} \underbrace{(p_B^\star + p_S^{\mathring{a}})}_{\substack{\text{Profit per}\\ \text{transaction}}}}_{\text{Payoff from working with } \mathring{a}}. \tag{2}$$

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

$$\frac{\delta}{1-\delta}\frac{\alpha}{1-\alpha}(q_B^\star + q_S^\star) \geq \frac{\alpha}{1-\alpha}(1-\kappa_S)\kappa_B(p_B^{\mathring{a}} + p_S^\star) + \frac{\alpha}{1-\alpha}(1-\kappa_B)\kappa_S(p_B^\star + p_S^{\mathring{a}})$$
$$\frac{\delta}{1-\delta}(v_B + v_S)(1-\kappa_S) \geq (1-\kappa_S)\kappa_B(p_B^{\mathring{a}} + p_S^\star) + (1-\kappa_B)\kappa_S(p_B^\star + p_S^{\mathring{a}});$$

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

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

[11] Note that $\frac{\alpha}{1-\alpha} = \frac{1}{|A|-1}$.

Hence, it is sufficient that

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

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

Intuitively, each agent a other than $\mathring{a}$ is unwilling to work with the 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 $\mathring{a}$ today.

Case 4: Mutual deviations. It is immediate that each agent's prescribed actions are optimal as they maximize profits this period and have no effect on future profits.

Making and Responding to Invitations in the Collusive Punishment Phase

The analysis of making and responding to invitations in the $\mathring{a}$-collusive punishment phase proceeds very similarly to the analysis of making and responding to invitations in the cooperation phase.

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

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

Case 3: Effective price deviation by $\mathring{a}$. 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 $q_B^\star$ and $q_S^\star$ instead of $p_B^\star$ and $p_S^\star$. In particular, the total payoff for a from following his prescribed actions is (cf. (1))

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

while the payoff for working with $\mathring{a}$ is given by (cf. (2))

$$\underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)(1-\kappa_B)(q_B^\star + q_S^\star)}_{\substack{\text{Profits from working with agents}\\ \text{other than } \mathring{a} \text{ this period}}} +$$

$$\underbrace{\underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_S)}_{\substack{\text{Mass of sellers}\\ \text{represented by } a}} \underbrace{\kappa_B}_{\substack{\text{Mass of buyers}\\ \text{represented by } \mathring{a}}} \underbrace{(p_B^{\mathring{a}} + q_S^\star)}_{\substack{\text{Profit per}\\ \text{transaction}}} + \underbrace{\frac{\alpha}{1-\alpha}(1-\kappa_B)}_{\substack{\text{Mass of buyers}\\ \text{represented by } a}} \underbrace{\kappa_S}_{\substack{\text{Mass of sellers}\\ \text{represented by } \mathring{a}}} \underbrace{(q_B^\star + p_S^{\mathring{a}})}_{\substack{\text{Profit per}\\ \text{transaction}}}}_{\text{Payoff from working with } \mathring{a}}.$$

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

Case 4: Ineffective Price Deviation by $\check{a} \neq \mathring{a}$. The analysis here is analogous to that of Case 2 during the cooperation phase.

Case 5: Effective price deviation by $\check{a} \neq \mathring{a}$. The analysis here is analogous to that of Case 3 in this subsection (where we considered an effective price deviation by $\mathring{a}$).

Case 6: Mutual deviations. It is immediate that each agent's prescribed actions are optimal as they maximize profits this period and have no effect on future profits.

Agent Selection in the Cooperation Phase

There are four 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 $p_B^\star \leq v_B$ to buyers, every agent offers a seller price $p_S^\star \leq v_S$ 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 $\mathring{a}$. If $\mathring{a}$ offers prices other than $(p_B^\star, p_S^\star)$, both buyers and sellers anticipate that $\mathring{a}$ will not have invitations accepted or accept any invitations. Thus, any buyer who is represented by $\mathring{a}$ will only see sellers represented by $\mathring{a}$ and, similarly, any seller who is represented by $\mathring{a}$ will only see buyers represented by $\mathring{a}$. 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 $\mathring{a}$ if and only if both

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

that is, if both

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

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

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

Case 4: Mutual deviations. It is immediate that these strategies are optimal for buyers and sellers given the strategies of agents.

Agent Selection in the Collusive Punishment Phase

The analysis of agent selection in the $\mathring{a}$-collusive punishment phase is analogous to that for agent selection in the cooperation phase. The analysis after either collusive pricing (Case 1) or mutual deviations (Case 6) is identical. After a effective or ineffective price deviation by either $\mathring{a}$ or another agent $\breve{a}$ (Cases 2–5) the analysis is similar to that of price deviations in the cooperation phase, except that the other agents are now offering prices of $q_B^\star$ and $q_S^\star$ instead of $p_B^\star$ and $p_S^\star$. Thus, we find that, for a price deviation to be effective during a collusive punishment phase, we need that

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

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°, q_S°). If an agent does post prices at or below (q_B°, q_S°), he obtains his full capacity of buyers and sellers but does not work with any other agents. Thus, his profits are given by

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

Since no agent 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 $\frac{1}{1-\delta}\alpha(p_B^\star + p_S^\star)$. 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

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

that is, so long as

$$\frac{\alpha}{1-\delta}(p_B^\star + p_S^\star) \geq \left(\left(v_B - \frac{1}{\kappa_S}(v_B - p_B^\star)\right) + \left(v_S - \frac{1}{\kappa_B}(v_S - p_S^\star)\right)\right)\kappa_B\kappa_S \qquad (3)$$

Hence, the highest sustainable profits are found by solving

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

subject to (3) and the individual rationality constraints for the buyers and sellers that $p_B^\star \leq v_B$ and $p_S^\star \leq v_S$. Solving this linear program yields

$$p_B^\star = \frac{(1-\delta)\kappa_B(v_B - \kappa_S(v_B + v_S)) + \alpha v_S}{(1-\delta)\kappa_B - \alpha}$$

$$p_S^\star = v_S.$$

Maximality of Surplus Extraction

Finally, we need to show that no surplus extraction greater than $p_B^\star + p_S^\star = p_B^\star + v_S$ can be maintained. If $p_B^\star$ is exactly v_B, then the individual rationality of buyers and sellers ensures that no greater surplus extraction is possible. If $p_B^\star$ is less than v_B, then any prices (p_B', p_S') that generate more surplus extraction than $(p_B^\star, p_S^\star)$ 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

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 can not 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 $(1-\delta)\kappa_B\kappa_S < (1-\delta)\kappa_S$, and continues to facilitate significant surplus extraction even as the number of agents grows large, and we might naïvely expect the market to become "perfectly competitive."

4 Policy Responses

4.1 Rebate Bans

Rebate bans, a feature of real estate regulation in several U.S. states, prohibit buyers' agents from offering "rebates," i.e., paying a buyer for the right to represent him in a real estate transaction. Han and Hong (2011) argue that rebate bans are anticompetitive and estimate a model where 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 $p^a_{B,t} \geq 0$ 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 such 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*

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

$$p^\star_S = v_S$$

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

Moreover, $\lim_{\alpha \to 0}(p^\star_B + p^\star_S) = v_B(1-\kappa_S) + v_S$, which is strictly higher than the highest sustainable industry profits without a rebate ban.

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 is negative; this happens when market concentration is sufficiently low, i.e., $\hat{\alpha} \leq \kappa_S \frac{v_S}{v_S+(1-\kappa_S)v_B}$. Thus, when market concentration is greater than $\kappa_S \frac{v_S}{v_S+(1-\kappa_S)v_B}$, the rebate ban has no effect. By contrast, when market concentration is less than $\kappa_S \frac{v_S}{v_S+(1-\kappa_S)v_B}$, the collusive buyer price $p^\star_B$ is set so

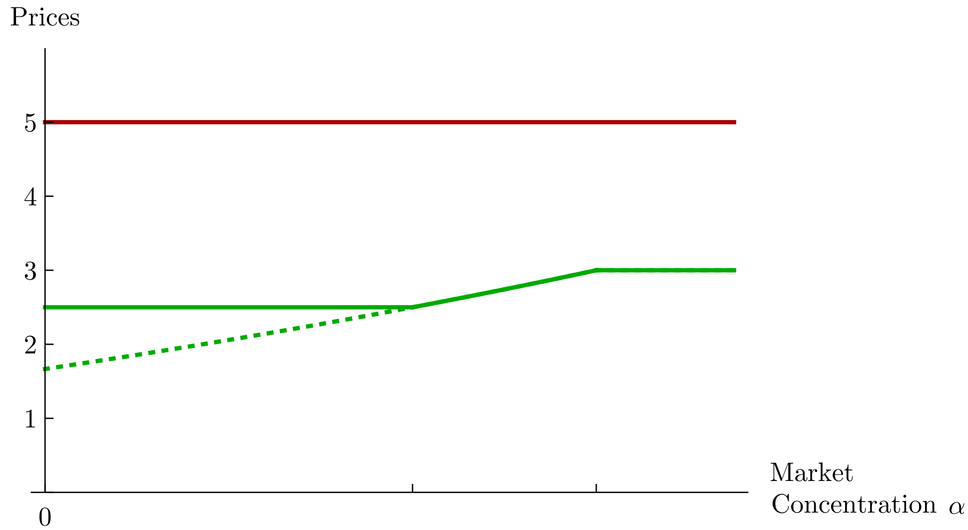

Figure 3: The prices supporting the highest sustainable profit equilibrium when a rebate ban is present. The dark red line is the seller price, $p_S^\star$, and the light green line is the buyer price, $p_B^\star$. The light green dashed line is the buyer price when no rebate ban is present; the seller 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 $\kappa_S = \frac{1}{6}$.

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 $p_B^\star$ (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 buyer agents and seller agents. In particular, while any seller agent can and likely does represent buyers at times, there are agents, and even firms of agents, who exclusively represent buyers.[12] Allowing for different types of agents may be important, as our model predicts that optimal collusion may require that agents subsidize buyers rather than charge buyers for access to the platform. Indeed, in the U.S. the price paid by a buyer to his agent is typically zero, despite some costs for the

[12]See, for example, the National Association of Exclusive Buyer Agents (`http://naeba.org/`).

agent for representing the buyer.

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 agent and a potentially high cost for acting as a buyer agent. *Buyer-exclusive agents* have a low cost for acting as a buyer agent and a very high cost for acting as a seller agent. This asymmetry 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. Agent heterogeneity can play an important role in the structure of optimal collusion. In particular, agents who only serve buyers 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 policies such as eliminating agency fees (which was suggested by Barwick (2018)).

4.2.1 Model

We modify the model by partitioning the set of agents A into the set of buyer-exclusive agents A_B and the 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. Consistent with the symmetric model, we normalize the cost of buyer-exclusive agents representing buyers and seller-proficient agents representing sellers to zero. Thus, the stage game payoff to a buyer-exclusive agent a is

$$|\mathbf{B}_t(a)| \sum_{\tilde{a} \in A^{a \Rightarrow}} \Big(|\mathbf{S}_t(\tilde{a})| (p^a_{B,t} + f^{a \leftarrow \tilde{a}}_t) \Big).$$

while the payoff to a seller-proficient agent is

$$|\mathbf{B}_t(a)| \sum_{\tilde{a}\in A^{a\Rightarrow}} \Big(|\mathbf{S}_t(\tilde{a})|(p^a_{B,t} + f^{a\leftarrow\tilde{a}}_t - c)\Big) + |\mathbf{S}_t(a)| \sum_{\tilde{a}\in A^{\Rightarrow a}} \Big(|\mathbf{B}_t(\tilde{a})|(p^a_{S,t} - f^{\tilde{a}\leftarrow a}_t)\Big).$$

4.2.2 Optimal Collusion

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

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

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

$$p^\star_S = v_S$$

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

Moreover, $\lim_{\sigma\to 0}(p^\star_B + p^\star_S) = (v_B + v_S)(1 - \kappa_S) + \kappa_S c$, 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

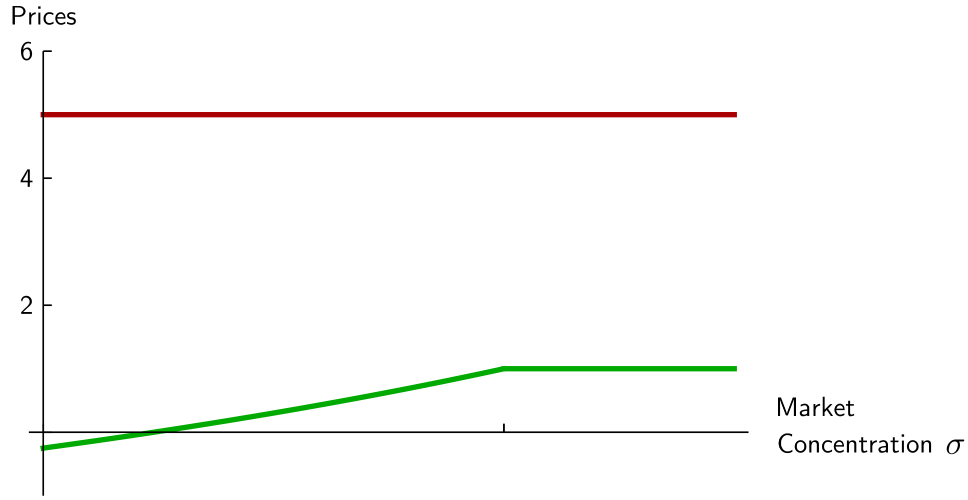

Figure 4: The prices supporting the highest sustainable profit equilibrium when seller-proficient agents work with buyers. The dark red line is the seller price, $p_S^\star$, and the blue line is the buyer price, $p_B^\star$. The light green line is the buyer price when agency fees are allowed; the seller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

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 $(q_B^\star, q_S^\star)$ and an agency fee $g^\star$. The second collusive punishment phase follows a price deviation in which the deviator attracts only sellers; this phase uses the same prices $(q_B^\star, q_S^\star)$ but a new agency fee $h^\star$. In both cases, the prices $(q_B^\star, q_S^\star)$ are given by

$$q_S^\star = v_S$$
$$q_B^\star = v_B - \kappa_S(v_B + v_S) + \kappa_S c.$$

Note that we do not need to worry about price deviations by either a buyer-exclusive or a seller-proficient agent $\mathring{a}$ on the buy-side; for such a deviation, each seller-proficient agent a (still) invites $\mathring{a}$ with a fee that obtains all of the surplus from a transaction between $\mathring{a}$'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 $(q_B^\star, q_S^\star)$, following a price deviation by some seller-proficient agent $\mathring{a}$ in the cooperation phase, no agent will work with $\mathring{a}$ this period. There are two relevant price deviations to consider: in the first, $\mathring{a}$ attracts both buyers and sellers, while in the second, $\mathring{a}$ just attracts sellers.[13]

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

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

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

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

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

[13] Note that it is easy to rule out the possibility that $\mathring{a}$ would attract only buyers by offering a price $p_B^{\mathring{a}} \leq p_B^\star$, by requiring other agents to invite $\mathring{a}$ with a fee of $c - p_B^{\mathring{a}}$, $\mathring{a}$ to accept any such invitation, and proceeding to Bertrand reversion in future periods. Thus, such a deviation by $\mathring{a}$ would obtain 0 profits this period; moreover, $\mathring{a}$ would obtain 0 profits in future periods.

that

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

Thus, as $q_B^\star + q_S^\star \geq (v_B + v_S)(1-\kappa_S)$, $p_B^\star + p_S^\star \leq v_B + v_S$, $\kappa_S \leq \kappa_B \leq \frac{1}{2}$, 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 $g^\star$ which satisfies both inequalities will be in $[-q_B^\star, q_S^\star]$.

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

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

Letting $h^\star = q_S^\star$, and recalling that $q_B^\star + q_S^\star \geq (v_B + v_S)(1-\kappa_S)$, $p_B^\star + p_S^\star \leq v_B + v_S$, 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 $q_B^\star$ and $q_S^\star$ 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 $(p_B^\star, p_S^\star)$ are given by

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

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

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

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

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

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

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

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 $g^\star \geq -q_B^\star$. We also require that a seller-proficient agent (other than $\mathring{a}$) will offer agency fees of $g^\star$; this is incentive-compatible so long as $g^\star \leq q_S^\star$.

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

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

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

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

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 $(q_B^\star, q_S^\star)$ 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

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

That is, we have chosen $(q_B^\star, q_S^\star)$ to provide the highest profits possible while ensuring that the punished agent can not 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

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

subject to

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

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 can not 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) in a coordination-proof subgame-perfect Nash equilibrium in which buyers only work with seller-proficient agents are achieved with prices*

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

$$p_S^\star = v_S$$

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

Moreover, industry revenue is weakly less than in the highest sustainable profit equilibrium with agency fees. However, $\lim_{\sigma \to 0}(p_B^\star + p_S^\star) = (v_B + v_S)(1 - \kappa_S) + \kappa_S c$, which is the same prices as under the highest sustainable profit equilibrium with agency fees.

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, the seller-proficient agents must be incentivized to represent buyers each period through dynamic considerations (when $p_B^\star$ is less that c, i.e., the price charged to buyers does not compensate the agent for the cost of representing that buyer). This 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 (where buyer-exclusive agents represent buyers and are compensated through agency fees). The equilibrium price for buyers is lower as individual

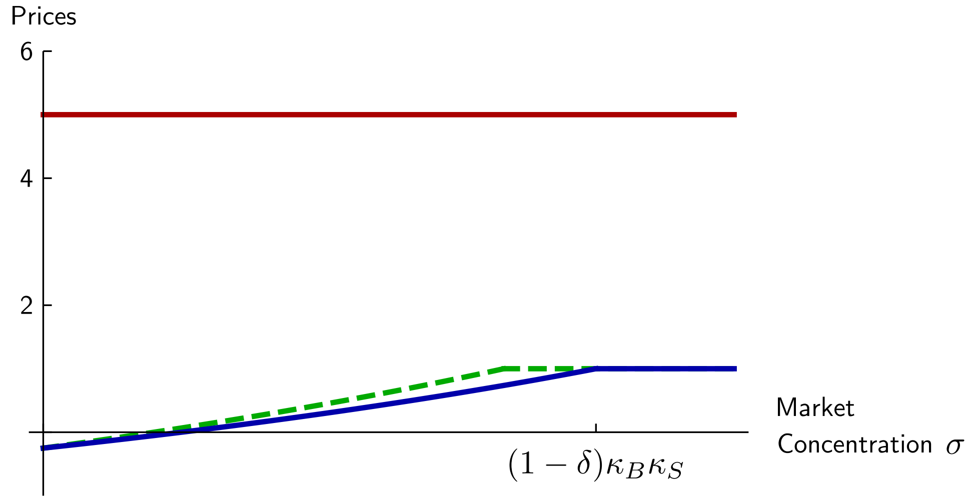

Figure 5: The prices supporting the highest sustainable profit equilibrium when seller-proficient agents work with buyers. The dark red line is the seller price, $p_S^\star$, and the blue line is the buyer price, $p_B^\star$. The dashed green line is the buyer price when agency fees are allowed; the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

equilibrium profits are lower (since seller-proficient agents now inefficiently 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) in a coordination-proof subgame-perfect Nash equilibrium*

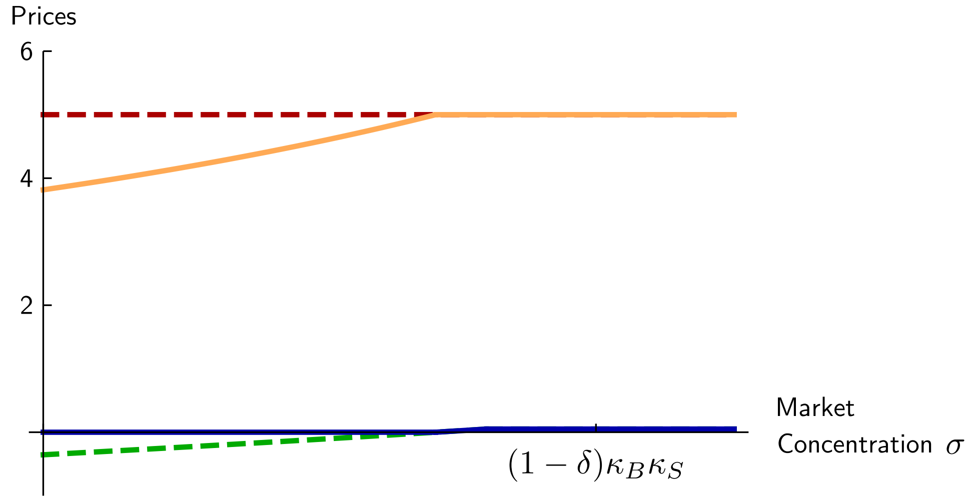

Figure 6: The prices supporting the highest sustainable profit equilibrium when seller-proficient agents work with buyers. The orange line is the seller price, $p_S^\star$, and the blue line is the buyer price, $p_B^\star$. The dashed green line is the buyer price when agency fees are allowed; the dashed red line is the seller price when agency fees are allowed. Here, $\delta = \frac{3}{4}$, $v_B = \frac{1}{20}$, $v_S = 5$, $\kappa_B = \frac{1}{3}$, $\kappa_S = \frac{1}{4}$, and $c = 1$.

in which buyers work with buyer-exclusive agents are achieved with prices

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

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

where $\hat{\sigma} = \frac{\sigma}{(1-\delta)\kappa_B}$.[14]

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 a charged a 0 price and sellers are charged a price a *less than* v_S. Figure 6

[14] Note that we can only have $\hat{\sigma} \leq \kappa_S(1-\frac{c}{v_S}-\frac{v_B}{v_S}\frac{1-\kappa_S}{\kappa_S})$ if the buyers' valuations are sufficiently low, i.e., $v_B \geq \frac{\kappa_S}{1-\kappa_S}(v_S - c)$.

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 $p_S^\star$, 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).[15]

Corollary 1. *If the value to buyers relative to sellers is sufficiently low, i.e.,* $v_B \frac{1-\kappa_S}{\kappa_S} + c < v_S$*, then* $\lim_{\sigma \to 0}(p_B^\star + p_S^\star) = v_B\left(\frac{\kappa_B}{\kappa_S} - 1\right) + v_S(1-\kappa_B) + \kappa_B c$*, which is less than the prices under the highest sustainable profit equilibrium with agency fees.*

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 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 show 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 (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.

[15] Note 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.

By formally modeling brokered intermediation, we can also understand the effects of proposed and existing policies. Here, we show that eliminating rebate bans, while not a panacea, reduces the scope for collusion. Similarly, eliminating agency fees also weakens 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.[16] 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 IPOs and syndicated lending (Hatfield et al., 2018; Cai et al., 2018). We also imagine these techniques may prove useful in analyzing business-to-business transactions and market entry.

[16] Repeated extensive form games have also been used to analyze vertical mergers (Nocke and White, 2007) and initial public offerings (Hatfield et al., 2018).

References

Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica*, 383–396. (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 4, 31, and 42.)

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 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.)

Cai, J., F. Eidam, A. Saunders, and S. Steffen (2018). Loan syndication structures and price collusion. SSRN #3250817. (Cited on page 43.)

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 page 2.)

Department of Justice (2005). Complaint. U.S. v. Kentucky Real Estate Commission. (Cited on pages 4 and 46.)

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 28.)

Hatfield, J. W., S. D. Kominers, R. Lowery, and J. M. Barry (2018). Collusion in markets with syndication. Working paper. (Cited on page 43.)

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 43.)

Levitt, S. D. and C. Syverson (2008). 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.)

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 3 and 13.)

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 42.)

Nocke, V. and L. White (2007). Do vertical mergers facilitate upstream collusion? *American Economic Review 97*(4), 1321–1339. (Cited on page 43.)

Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review 100*(4), 1642–72. (Cited on page 8.)

White, A. and E. G. Weyl (2016). Insulated platform competition. Working paper, Tsinghua University. (Cited on page 8.)

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 broker 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

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

where $a \succ_b \hat{\alpha}$ if a is ranked higher than $\hat{\alpha}$ 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

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. □

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 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 Proof of Theorem 2

To be written.

C.3 Proof of Theorem 3

To be written.

C.4 Proof of Theorem 4

To be written.

C.5 Proof of Theorem 5

To be written.

D Analyzing Asymmetric Equilibria in the Stage Game

To be written.