

Can forgetful sellers be better off?
Impact of information in an ultimatum
price-setting game with learning

Christopher Cotton*

Cornell University

October 14, 2005

Abstract

This paper introduces learning dynamics into a posted-offer pricing game, in which sellers observe past-period transactions before announcing a take-it or leave-it price, and buyers either accept or reject the announced price. We consider the impact that seller access to information regarding past transaction has on the long-term prices, and show that when sellers have imperfect information about the past, the long-term average sale price may be higher than when sellers perfectly observe the entire history of the game. It follows that limiting seller information can improve their long-term average welfare, and total long-term average sales revenue. This has interesting implications regarding firm incentives to provide information to their managers and sales agents.

*I appreciate the comments of Talia Bar, Larry Blume, Drew Fudenberg, Ani Guerdjikova, Sjaak Hurkens, Ted O'Donoghue, Mike Waldman, attendants at the International Game Theory Conference at Stony Brook, and seminar participants at Cornell University. All mistakes remain my own. Contact information: Department of Economics, Cornell University, Ithaca, NY; csc35@cornell.edu. JEL Classification Numbers: C73, D83. Keywords: learning, posted price games, limited memory, management incentives.

1 Introduction

Basic models of market interaction assume the price-setting sellers have full knowledge of the price-taking buyers' willingness to pay for the product, and are thereby able to accurately set a price that maximizes profits. In reality, however, this is often not the case. With few exceptions, sellers are not endowed with a perfect understanding of buyer demand. Additionally, undertaking costly preliminary market research rarely results in certainty regarding the profit maximizing price, although it may narrow down a potential price range. Instead, sellers usually learn about buyer willingness to pay over time, as they observe past transactions and update their beliefs as more information becomes available.

To deal with similar dynamics, a number of learning models have developed in the literature (for an overview, see Weibull 1995, or Young 2005). In these models, the long-term outcome depends on the learning rules, or more specifically, on the agent sophistication and ability to observe past periods, which may also be referred to as their memory. Processes in which agents observe only the outcome of the previous period will likely have a different long-term outcome than processes in which agents observe, and therefore learn from, the entire history of the game. For example, Hurkens (1995) shows how large enough memory of recent periods can limit the potential strategy set of the agents in a learning game. Young (1993b) shows how sufficiently small memory can increase the likelihood of certain outcomes in a bargaining game, such as 50-50 division. This paper incorporates learning dynamics into a simple ultimatum posted price game, and considers how different memory limits can influence the long-run prices that are announced over the course of the game. Specifically, we consider whether limiting seller memory (or ability to observe past periods) can result in higher average sale prices, thereby increasing long-run seller welfare, or whether sellers are strictly better off if they can observe all past transactions.

This paper's model is built on the framework used in Young's evolutionary bargaining model (1993b). Similar to Young's model, this paper considers an infinitely repeated stage game, where, in each period, two agents are randomly matched to play against each other.

In Young's model, these agents were a landlord and a potential tenant, and in our model, they are a seller and a potential buyer of a good or service. Additionally, rather than the split-the-dollar bargaining game employed by Young, we use a repeated posted offer game in which only the seller is able to announce a take-it or leave-it price, and the buyer can only accept or reject this price. These differences allow us to address a typical market transaction in which one party sets a non-negotiable price for a good or service.¹

In each period, an agent from the pool of sellers is randomly matched with a buyer to play that period's stage game. After being selected, the seller observes a partial or complete history of the game, including prices announced in past periods and whether they were accepted or rejected. Then the seller updates her expectations regarding the buyer's valuation, and announces her own price. The buyer can either accept or reject the price. If he accepts the price, the buyer receives a payoff equal to the difference between his valuation and the price, and the seller receives payoff equal to the price. If he rejects the price, both the buyer and the seller receive a payoff of zero. Even though no negotiation takes place between the seller and the buyer in this model, the interpretation of this model need not be so limited. In situations where a sales agent and a buyer expect to negotiate over a price, this model may still apply so long as the range of prices over which they negotiate is relatively small compared to the entire range of possible buyer valuations.

In this model's framework, we analyze three cases of seller information. In the first case, sellers always observe all past transactions, and are therefore perfectly informed of the entire history of the game. In contrast, the second and third cases both involve imperfect seller information. In the second case, sellers observe only the m most recent transactions. In the third case, sellers observe each past period transaction with an independent positive probability, and therefore it is possible to observe any subset of the complete history of the game.

¹In contrast, in Young's (1993b) primary example, randomly matched landlords and tenants both announce demands over the division of a crop. They both receive their announced division if and only if the sum of the two divisions is less than or equal to the total.

For all three of these cases, we show that the pricing process settles to long-run *conventions*, which we define in detail in the following section. Once the process settles to a long-run convention, the price associated with the convention is the only price at which sales take place in any future period. Convention prices are necessarily no larger than the buyer valuation, and are therefore always accepted. When sellers have perfect information of past transactions (case 1), the process settles to a permanent convention, p_{pc}^* . Once established, a permanent convention price is announced by sellers in all future periods. When observations are limited to the most recent periods (case 2), the process eventually settles to a repeatedly-reoccurring convention, p_{rrc}^* . After p_{rrc}^* is established, sellers occasionally experiment with higher prices that are rejected by the buyer; however, the process always returns to p_{rrc}^* . In the long run, the probability that p_{rrc}^* is announced and therefore accepted in any period approaches some constant. When sellers observe each past period with positive probability (case 3), the process achieves a convention in probability, p_{cip}^* . A process achieves a convention in probability p_{cip}^* when the probability that any seller announces p_{cip}^* approaches 1 in the long run.

When a seller announces her price in any period, the price is determined after considering the observed past accepted and rejected prices. A seller will always announce a price at least as large as the maximum past accepted price she observed, p^{\min} , and strictly less than the minimum past rejected price she observed, p^{\max} . This is because she knows that the buyer valuation is within this *potential price range*. When a convention is achieved, all sellers prefer to announce the convention price $p^* = p^{\min}$ rather than any other $p \in [p^{\min}, p^{\max})$.

Under reasonable conditions, providing sellers with imperfect observation of past transactions (whether by limiting memory to the most recent periods or randomly determining which past periods the sellers observe) can cause the process to achieve a long-term convention price that is closer to the buyer valuation than the long-term convention price that would be established if sellers perfectly observed all past transactions. The reasons and requirements for this are presented in detail in the body of the paper. In summary, the

primary reason that imperfect information can result in a higher long-term price compared to perfect information is that it may result in sellers observing potential price ranges that are not observed under perfect information. For example, suppose that sellers observe only the m most recent transactions, and the process converges to a temporary convention with potential price range $[p_{pc}^*, p^{\max})$, where $p_t = p^{\max}$ was announced (and rejected) in period t . From when this convention is achieved through period $t + m$ all sellers announce p_{pc}^* , and the potential price range remains constant between periods. With perfect information, we would remain here indefinitely. However, with m -period memory, the seller in period $t + m + 1$ no longer observes p^{\max} since it was originally announced more than m periods before. When this happens, the upper bound of the seller's potential price range reverts back to the absolute maximum price that the seller believes possible, and the period $t + m + 1$ seller now chooses a price out of a larger potential range. The seller may now choose a different price rather than p_{pc}^* . This new experimentation can result in a new price being announced between p_{pc}^* and the buyer valuation. If this happens, any future convention price, including p_{rrc}^* , will be at least as large as this new price.

Alternatively, when sellers observe each past period with positive probability, there is a positive probability that a seller observe any subset of the complete game history. Now, suppose that at some point a seller observes a selection of past price histories such that she announces price p_{pc}^* . In the periods that follow, there is a positive probability that sellers observe price histories such that $p^{\min} = p_{pc}^*$, and also a positive probability that p^{\max} equals any of the past rejected prices. Therefore, when sellers observe $p^{\min} = p_{pc}^*$, they do not necessarily observe the same p^{\max} they would have under perfect information (for which p_{pc}^* is the optimal price for all sellers). Instead, their observation set may result in a p^{\max} such that the resulting potential price range $[p_{pc}^*, p^{\max})$ was never observed as the process converged under perfect information. When this happens, p_{pc}^* may not be the optimal price announcement, and the seller may announce some new price larger than p_{pc}^* and less than the buyer valuation. If this happens, any future convention price, including p_{cip}^* , will be at

least as large as this new price.

After we analyze the three cases of information, we consider an application of the model that treats the entire pool of sellers as a firm, and the individual sales agents within the pool as employees or managers acting on behalf of the firm. Because sellers' incentives often differ from those of the firm owners, sellers may choose actions that do not maximize firm earnings (Basu, et al 1985; Holmstrom 1999). Our model can address the firm-employee relationship when firm owners are more concerned with long-term performance, than performance on individual projects or sales. For example, consider an automobile dealership that is primarily concerned about annual or quarterly sales revenue, while its sales agents may be primarily concerned with immediate sales. To help align employee behavior with firm preferences, the firms may benefit from encouraging risk taking among its sales agents, since risk taking can push long-term convention prices closer to the buyer valuation. This result supports a variety of past literature that shows how a firm can benefit from affecting employee risk preferences through compensation schemes such as standard employment contracts, sales contests, quotas, and promotion rules (see for example Wilson 1968; Ross 1973; and Gaba and Kalra 1999). Our analysis contributes to this literature by showing how firms may also achieve similar results by limiting seller access to information regarding past transactions. Given costs associated with the standard compensation tools, information limits may be a less costly or more effective alternative. As far as we know, this is the first paper showing how employee information limits can have equivalent result on employee behavior as more standard compensation tools.

Section 2 describes the model and defines the dynamic price process. Section 3 analyses the convergence of the process under the three cases of information limits and considers the requirements for imperfect information to result in higher average seller utility compared with perfect information of past transactions. The section also considers the application of the model to the firm-seller relationship, and discusses what happens when the sellers do not fully discount future periods. The final section of the paper presents the concluding remarks,

including a brief discussion regarding possible expansions of this paper.

2 Model

There are two pools of individuals, A and B . A represents the class of sellers of a good or service, and B represents the class of buyers of that good or service. The game has an infinite time horizon, and in each period $t = 1, 2, 3, \dots$ one seller α is drawn at random from class A , and one buyer β is drawn at random from class B . Where required for clarity, we use the notation α_t to represent the seller that is selected to play in period t . α then observes some past period transactions, and can update her expectations regarding the buyer's willingness to pay for the good. The seller then announces a price p for the good or service, and β chooses whether to accept or reject the price. If the buyer, β , accepts the price, he receives a payoff equal to the difference between his valuation and the price ($V_\beta - p$), and the seller receives a payoff equal to the price p . If the buyer rejects the price, the transaction does not take place, and both the buyer and seller receive nothing. This price setting process changes the underlying game in Young's bargaining model (1993b), but retains many similarities in structure with his model.

Buyers are defined by their valuation for the good, which is denoted V_β . For technical reasons, similar to Young (1993b), we assume that there exists a finite number of feasible prices and valuations. This assures that the process can converge to a price convention in which the same prices are necessarily announced in sequential periods, rather than only allowing the process to converge to prices arbitrarily close to the convention value.² Let there be r feasible prices and valuations along the continuum $[V_{\min}, V_{\max}]$, and let δ denote the precision of the price range, where $\delta = \frac{V_{\max} - V_{\min}}{r-1}$. Therefore, p and V_β are in the set $\{V_{\min}, V_{\min} + \delta, \dots, V_{\max} - \delta, V_{\max}\}$; however, through some abuse of notation, we refer to this

²Most results do not change if we assume a continuum of potential prices and allow a convention at price p to be achieved when price announcements are necessarily within a neighborhood sufficiently close to p . Further development of the convention concept can be found in Young (1993a).

discrete set of potential prices by the notation $D = [V_{\min}, V_{\max}]$.³

Seller types can differ in terms of their utility functions, and their ex ante expectations regarding potential buyer valuations. Seller utility depends on the price received for the good. Generally, sellers may have different utility functions $u_\alpha(p)$; although, for all sellers, $u(V_{\min}) > 0$, $u(0) = 0$, and $u'(p) > 0$. If the individual sellers are thought of as risk-neutral firms within a trade organization, then $u''(p) = 0$. If the sellers are sales agents within a firm, as generally assumed, then the structure of their utility functions are likely to depend on the compensation agreements they have with their firm (Basu, et al 1985; Holmstrom 1999). Therefore, $u''(p) > 0$, $u''(p) = 0$, and $u''(p) < 0$ are all possible.

Additionally, sellers cannot observe buyer willingness to pay for the good. However, sellers do have expectations regarding the distribution of the buyer valuation. Let the CDF $\bar{F}_\alpha(\cdot)$ represent seller α 's ex ante beliefs regarding the possible distribution of the buyer valuation V_β , and $\bar{f}_\alpha(\cdot)$ represent the distribution's density, such that $\bar{f}_\alpha(P) > 0$ for all $P \in [V_{\min}, V_{\max}]$. Therefore, when selected to play the game, prior to updating beliefs, sellers believe that any of the potential valuations on D are possible.

Because the primary focus of this paper involves seller decision process and the impact of their actions on seller and firm welfare, a few simplifying assumptions are made regarding the class of buyers. First, a buyer's type is defined by its valuation V_β alone. Second, the class of buyers is homogeneous in valuation. This second assumption may also be thought of as sellers having the ability to distinguish buyers of different types, even if they cannot observe the valuation associated with the types.⁴ Finally, buyers are non-strategic, and accept any price less than their valuation, $p \leq V_\beta$. This may be because buyers completely discount future periods, or because they are selected to play the ultimatum game at most once, then are replaced by an identical buyer.

³ V_{\min} and V_{\max} are defined by the sellers' beliefs. We assume that the sellers are correct in that $V_\beta \in [V_{\min}, V_{\max}]$.

⁴For example, an auto mechanic may not know the valuation that each driver places on his services. However, if all sellers that drive the same type of car have the same valuation, and he observes car type, then the results in this model continue to hold.

On the other hand, the pool of sellers may be composed of different seller types, as defined by their utility functions and prior beliefs regarding buyer valuation. We can interpret the pool of sellers in a variety of ways. In the first interpretation, sellers are non-strategic, either because they completely discount future periods, or, similar to Young (1993b), they only play the ultimatum game once, then exit the pool and are replaced by an identical agent. In the second interpretation, there are a finite number of sellers, each of which has a positive probability of being drawn to play in any given period. In this case, the finite number of sellers can be thought of as a sales or management team within the firm. Their concern regarding future periods depends on their individual discount rates, as well as the probability they will be drawn in any given period. In the third interpretation, there is a single seller (executive, manager, etc.) who is selected to play the game in each period for sure. In this case, the seller's concern for the future depends only on her discount rate. This paper and its analyses focus on the case where sellers are non-strategic regarding future periods. However, under reasonable conditions, the conclusions in this paper can be generalized for any of the three interpretations above. These conditions are further discussed below.

Let p_t denote the price announced by the seller in period t , and let $a_t \in \{0, 1\}$ be an indicator variable describing whether p_t was accepted by the buyer in that period. $a_t = 1$ if and only if the buyer accepted p_t . Therefore, the set (p_t, a_t) denotes the *price history* for period t , and the sequence $H_t = \{(p_1, a_1), (p_2, a_2), \dots, (p_{t-1}, a_{t-1})\}$ denotes the *complete price history* of the game up to period t . For consistency, let H_t^A denote the set of all past prices within H_t that were accepted, and H_t^R denote the set of all past prices within H_t that were rejected. Additionally, K_t denotes the set of price histories observed by the seller selected to play the ultimatum game in period t . The analysis begins by considering the case where $K_t = H_t$, then considers the impact of limiting seller information such that $K_t \subset H_t$.

When the class of buyers is homogeneous, they share the same valuation of the good. Therefore, the seller knows that the buyer valuation is at least as large as the maximum accepted price and less than the minimum rejected price in the seller's observed price history.

To formalize this process, let $P_A \subset K_t$ be the subset of all observed past-period prices that were accepted; and let $P_R \subset K_t$ be the subset of observed past-period prices that were rejected.⁵ Then we can define $p_t^{\min} = \max \{p \mid p \in P_A\}$ and $p_t^{\max} = \min \{p \mid p \in P_R\}$. Therefore, p_t^{\min} is the maximum observed accepted price and p_t^{\max} is the minimum observed rejected price at period t . If P_A is an empty set, then $p_t^{\min} = V_{\min}$. If P_R is an empty set, then $p_t^{\max} = \delta + V_{\max}$. The set of feasible prices and valuations $\{p^{\min}, p^{\min} + \delta, \dots, p^{\max} - \delta\}$ defines the *potential price range*. Through abuse of notation, this discrete set is referred to by the notation $[p^{\min}, p^{\max}]$. Therefore, the relevant data contained in the observation set is contained in the values of p^{\min} and p^{\max} , and the seller effectively ignores the rest of the observations.

After observing K_t , the seller then updates her beliefs regarding the buyer valuation, forming a new CDF $F_{\alpha_t}(P \mid K_t)$ defined by the density function $f_{\alpha_t}(P \mid K_t)$. With a homogeneous pool of buyers, these functions can alternatively be written $F_{\alpha_t}(P \mid p_t^{\min}, p_t^{\max})$ and $f_{\alpha_t}(P \mid p_t^{\min}, p_t^{\max})$. When a seller updates its priors, given its observations of past period transactions:

$$f_{\alpha_t}(P \mid p_t^{\min}, p_t^{\max}) = \frac{\bar{f}_{\alpha}(P)}{\bar{F}_{\alpha}(p^{\max} - \delta) - \bar{F}_{\alpha}(p^{\min} - \delta)} \quad (1)$$

for all $P \in [p^{\min}, p^{\max}]$, and 0 otherwise. Sellers update their beliefs in the same way whether they observe a partial or a complete history of the game.⁶

⁵The A in P_A is not related to the class of firms, also defined by A .

⁶In some sense, this means that the sellers do not fully understand the multiple-period game in which they play a role. When they observe a partial history of the game, they do not use their information to infer what the complete history of the game may look like. An alternative assumption would give the individual sellers complete understanding the multiple-period game, and enable them to determine the exact probability of all possible true state of the game, given their observations, and act accordingly. With these fully-rational sellers, the analysis produces similar results regarding the potential benefits of limited memory. However, the range of parameter values over which the benefits hold are reduced. We believe that the assumptions involving the agents in this paper's model better represent the participants in most real world transactions compared to fully-rational agents.

At this point, I do not provide much of a discussion regarding the model if the class of buyers is heterogeneous. However, I will briefly discuss the differences in setup, between the cases of heterogeneous and homogeneous buyers. In the case of heterogeneous buyers, the seller can no longer update its CDF while only considering the highest accepted price demand, and the lowest rejected price demand, as can be done for homogeneous sellers. Instead, the sellers must allow for different buyers to have different valuations. This may be done through a general updated CDF $F(P \mid K_t)$ which is defined by the densities $f(P \mid K_t)$, such

Throughout this paper, we concentrate on Nash equilibria in which a buyer always accepts a price when it is less than or equal to the buyer's valuation, where $p \leq V_\beta$. The seller chooses p , and gets p if and only if $p \leq V_\beta$. The probability that a seller believes that a price p will be accepted is therefore given by an expression involving the updated CDF $F_\alpha(\cdot)$:

$$\Pr \{p \leq V_\beta\} = 1 - F_{\alpha_t}(p - \delta \mid K_t) \quad (2)$$

$$\Pr \{p \leq V_\beta\} = 1 - F_{\alpha_t}(p - \delta \mid p_t^{\min}, p_t^{\max}) \quad (3)$$

Including δ in the expression is necessary given the properties of the discrete potential price set, where $\Pr \{p < V_\beta\} = 1 - F(p)$, and δ is the minimum possible increase in price.

Therefore, seller α solves:

$$\max_{p \in D} u_\alpha(p) [1 - F_{\alpha_t}(p - \delta \mid K_t)] \quad (4)$$

The agents' response rules determine a stationary Markov chain. Let $\theta_\alpha(p \mid H_t)$ be the conditional probability that seller α announces price p given that α is selected to play the game at time t , and that the history of the game is given by H_t . Assume that θ_α is a *best reply distribution*; that is, $\theta_\alpha(p \mid H_t) > 0$ if and only if $p \in \arg \max_p u_\alpha(p) [1 - F_{\alpha_t}(p - \delta \mid K_t)]$. H_{t+1} is a *successor* of H_t if $H_t \subset H_{t+1}$, such that H_{t+1} has the same price history as H_t up through time $t - 1$, but also has an additional price history for period t given by (p_t, a_t) . Let $\pi(\alpha)$ be the probability that α is drawn to play the period ultimatum game in any period. Every $\alpha \in A$ has a positive probability of being drawn, though it is not necessarily the same probability for all agents. If the process has history H_t at time t , then it has history H_{t+1}

that

$$f(P \mid K_t) = \frac{1}{k} \sum_{n \in K} \left[\frac{\bar{f}(P)}{1 - \bar{F}(p_n)} a_n b_n + \frac{\bar{f}(P)}{\bar{F}(p_n)} (1 - a_n) (1 - b_n) \right]$$

where K is the set describing the k past periods observed by α , and $b_n \in \{0, 1\}$ is an indicator variable such that $b_n = 1$ iff $P \leq p_n$. $F(P)$ may reasonably take another form, so long as it combines the seller's prior beliefs and observations into a new CDF.

at time $t + 1$ with probability

$$\Phi_{H_t H_{t+1}} = \sum_{\alpha \in A} \pi(\alpha) \theta_\alpha(p | H_t) \quad (5)$$

If H_{t+1} is not a successor of H_t , then $\Phi_{H_t H_{t+1}} = 0$. This Markov process will be called the *dynamic price process* with precision δ , memory m , and best reply distribution p_α .⁷

3 Analysis

The analysis begins by introducing the concepts that the analysis incorporates. We then consider the case in which sellers learn about the complete history of the game before announcing a price in any period. We can think of this as the sellers being perfectly informed about all past sales, or that they observe all past transactions and have infinite length memory. After analyzing the case of perfect information, we consider the impact that limiting seller access to information (or limiting memory) has on the long-term convergence of the price-setting process. We consider two cases of limited memory. In the first, sellers observe only the most recent m period transactions. In the second, the sellers observe each period transaction with probability $\rho \in (0, 1)$. Following these analyses, we consider an application of the model to address the relationship between firms and their sales agents, and discuss how our results remain unchanged when sellers do not completely discount future periods.

3.1 Conventions

In any period t , after observing a potential price range $[p_t^{\min}, p_t^{\max}]$, the randomly drawn seller chooses a price p_t to maximize her expected utility. She solves:

$$\max_{p \in D} u_{\alpha_t}(p) [1 - F_{\alpha_t}(p - \delta | p_t^{\min}, p_t^{\max})] \quad (6)$$

⁷If the pool of buyers is composed of different types, then the dynamic price demand process evolves according to $\Phi_{H_t H_{t+1}} = \sum_{\alpha \in A} \sum_{\beta \in B} \pi(\alpha, \beta) \theta_\alpha(p | H_t)$, where $\pi(\alpha, \beta)$ is the probability that α and β are drawn to play against each other in any period. $\pi(\alpha) = \sum_{\beta} \pi(\alpha, \beta)$.

A seller always selects a price contained within her potential price range, since any price above the range is rejected for sure, and any price below the range always results in lower earnings compared to p_t^{\min} , which is accepted for sure. Additionally, note that p_t^{\min} is the only price in the potential price range that is accepted for sure, and for every other price within the range, the seller recognizes a positive probability that the buyer will reject the price. A seller is said to *experiment* with price if she announces a price that may be accepted or rejected both with positive probability, which is true for all prices within (p_t^{\min}, p_t^{\max}) . Alternatively, the seller can not experiment and choose to receive price p_t^{\min} for sure.

The process is said to *converge* in period t if the potential price range in period $t + 1$ is a non-equal subset of the potential price range in period t . Under perfect information of game history, the potential price range can only converge or remain constant, and converges in period t if and only if the period- t seller experiments and announces $p_t \in (p_t^{\min}, p_t^{\max})$. Under imperfect information of game history, the potential price range can potentially expand between periods, and convergence is not guaranteed.

The analysis is primarily concerned with comparing the long-term (asymptotic) prices established given different rules regarding the ability of the class of sellers to observe past period transactions. To accomplish this, we focus on the establishment of conventions, or states in which all possible sellers in any period prefer to announce the lower bound of the potential price range, which they get for sure, rather than any other possible price. When a convention is established, all sellers believe that the potential benefit (in terms of increased sales price) from experimenting, is out weighed by the potential loss due to the possibility that the price is rejected. The basic concept of a convention is formalized in a pair of related definitions.⁸

Definition 1 *The process achieves a **common action** at time t and price p^* when $p^* \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | K_t)]$ for all $\alpha \in A$.*

⁸However, although the definitions of conventions are similar between this paper and Young (1993b), the differences in the models' framework mean that the concept must be redefined here. Young defines a convention: "A state \mathbf{s} is a *convention* if it consists of some fixed division $(x, 1 - x)$ repeated m times in succession, where $x \in D$ and $0 < x < 1$."

Definition 2 *The process maintains a **convention** at time $t + 1$ when the process achieves a common action p^* at time t , and $p^* \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p_{t+1}^{\min}, p_{t+1}^{\max})]$ implies that $p^* \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p_{t+1}^{\min}, p_{t+1}^{\max})]$ with probability one.*

If we assume that any $p \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p^{\min}, p^{\max})]$ is possible, then we can be assured that the p^* associated with a convention in period t is the unique solution to all seller's maximization problems $\max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p^{\min}, p^{\max})]$. It also follows that $p_t^* = p_t^{\min}$ when there is a convention in period t . In the analysis that follows, we are concerned with three specific refinements of this concept that arise in the long-run. The first two refinements, permanent and repeatedly-reoccurring conventions, strengthen the standard definition of a convention. A *permanent convention* refers to a convention that, once established, remains established in all future periods of the game. As discussed in detail below, this concept necessarily results in the case of perfect information of past transactions.

Definition 3 *A **permanent convention** is achieved at time t and price p^* when $p^* \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | K_t)]$ for all possible α , and $p_t = p^*$ implies $p_s = p^*$ for all $s > t$.*

Generally, once a convention is established, the process can leave the convention, and possibly converges to a new convention with a different potential price range than the original convention. A special case of this temporary convention happens when the process necessarily converges back to the same convention that it left. This concept is called a *repeatedly-reoccurring convention*. Once a repeatedly-reoccurring convention is established, it is the only convention that is achieved through all future periods of the game; however, the process routinely leaves the convention, only to return to it again in a future period.

Definition 4 *p^* is a **repeatedly-reoccurring convention** if there exists a t such that in all periods following period t , the process fluctuates between being in convention p^* and not being in any convention at all.*

The third concept weakens the standard definition of a convention. Unlike after one of the other types of conventions results, when a *convention in probability* is established, the probability that the convention price is announced in any given period may be less than one. However, over the long run, this probability approaches one.

Definition 5 p^* is a **convention in probability** if as $t \rightarrow \infty$, $\Pr(p_t = p^*) \rightarrow 1$ (the probability that $p_t = p^*$ approaches 1), for all sellers.

A permanent convention is a special case of a convention in probability.

3.2 Perfect Information of Past Transactions

When sellers have perfect information of past transactions, they observe all of the past prices in the game, and whether the prices were accepted or rejected by the buyers. This means that the set of prices observed by a seller always contains the prices observed by sellers in past periods, and may also include some additional more recent prices. Since relevant seller information is completely defined by the potential price range, it follows that $[p_s^{\min}, p_s^{\max}] \subset [p_t^{\min}, p_t^{\max}]$ for all periods t and $s > t$, with equality possible. Therefore, under perfect information of past transactions, the potential price range either converges or remains constant between periods, and never expands.

If $p_t \in (p_t^{\min}, p_t^{\max})$, then the seller announces a price that has not been announced before, and all future sellers can observe whether this new price was accepted or rejected. If p_t is rejected, it becomes the new minimum rejected price observed by the seller in period $t + 1$; if it is accepted, it becomes the new maximum accepted price observed by the seller in period $t + 1$. Experimentation necessarily results in the potential price range in period $t + 1$ being a non-equal subset of the potential price range in period t . Therefore, under perfect information of past transactions, experimentation necessarily results in convergence of the potential price range.

Alternatively, if the seller chooses $p_t = p_t^{\min}$, then no new price announcements are available to the seller in the following period, and the process does not converge in period t . Because sellers may differ in terms of their expectations and risk preferences, α_t 's announcement of $p_t = p_t^{\min}$ does not imply that all other possible sellers also find it optimal to select p_t^{\min} . If there does exist a seller $\alpha \in A$ that prefers some $p \in (p_t^{\min}, p_t^{\max})$ to p_t^{\min} at time t , then given a infinite time horizon, such a seller will eventually be drawn, and the process will necessarily converge further. Only when, given the potential price range, all possible sellers prefer to announce the lower bound on the potential price range p_t^{\min} will the process no longer converge. When this happens, the process achieves a *convention*. Because there only exists a finite number of possible prices, the process can only converge a finite number of times, and in the long run a convention will eventually be achieved. Additionally, when a convention is established, no new price observations become available to sellers in the following period; and perfect information implies that all price information that was available to a seller in one period is also observed by all sellers in future periods. These two factors mean that once a convention is established, all future periods result in an equivalent potential price range, and therefore the convention is a permanent convention.

Proposition 6 *When sellers have perfect knowledge of the game history, the process almost surely converges to a permanent convention.*

This proposition tells us that under perfect information of game history the process will eventually settle to a state in which the same price is announced and accepted period after period, and that the process remains in that state indefinitely. As the following proposition states, the permanent convention price may be less than the buyers' valuation.

Proposition 7 *If $V_\beta > V_{\min}$, then, given agents' utility functions, there exists some distribution of prior beliefs regarding V_β such that the permanent convention price p^* is less than the buyers' valuation ($p^* < V_\beta$).⁹*

⁹Alternatively, this proposition can be changed to fix the agents' ex ante beliefs, and then increase the risk aversion of sellers until a similar result is established.

This means that the long-term price may be less than buyers are willing to pay, even if the sellers have perfect information regarding the transaction history of the game. This result follows simply from the sellers' expected utility maximization problem. In any period, sellers can always announce the minimum potential price, and receive it for sure. Although experimenting with a higher price results in the possibility of receiving a higher price, it also results in the increased possibility of having their price rejected and receiving nothing. This result may hold, even if the sellers are highly risk seeking in sales price.

Consider this simple example. Let all members of the class of sellers be risk neutral in sales price, and initially believe that the possible buyer valuation be uniformly distributed from $V_{\min} = 1$, and $V_{\max} = 100$, where $\delta = 1$, and actual buyer valuation $V_\beta = 45$. The first seller to play the game observes $p_1^{\min} = 1$ and $p_1^{\max} = 101$, and chooses $p_1 = 50$ to maximize her expected utility, which is rejected. In the following period, $p_2^{\min} = 1$ and $p_2^{\max} = 50$, and α_2 announces $p_2 = 25$, which is accepted. Now, in the third period, $p_3^{\min} = 25$ and $p_3^{\max} = 50$, and α_3 elects not to experiment, and announces $p_3 = 25$, which she receives for sure. No further price information is available to the seller in the following period, or any future period, and therefore the sellers in the following periods choose the same price $p_t = 25$ for all $t = 3, 4, \dots$. Therefore, a permanent convention is achieved at a price significantly less than the buyer valuation of 45.

Even when sellers can observe all past transactions in the game, the long-term price may be less than the buyers are willing to pay. Therefore, p_{pc}^* may be less than the buyer valuation V_β . The difference between the permanent convention price p_{pc}^* and V_β depends on the evolution of the price process as it converges to a convention. The path of convergence depends on the ex ante expectations that sellers have over the buyer's valuations $\bar{F}_\alpha(\cdot)$, the form of the sellers' utility functions $u_\alpha(\cdot)$, and the random selection of seller types to play each stage game. This is true for the convergence under perfect information of past periods, as well as under bounded memory, which we analyze in the following sections. If sellers place low enough probability on the buyer valuation being equal to the actual valuation, then even

the extremely risk-seeking seller may choose a lower price. Alternatively, even if sellers have highly-accurate beliefs, if they are also highly risk averse, the permanent convention price can still be less than the buyer valuation.

Typically, we treat the seller risk preferences and ex ante beliefs regarding the buyer valuation as fixed. However, it is interesting to consider what would happen to the long-term convention price if we altered these factors. Remember, we restrict the priors of the sellers by requiring them all to place a positive ex ante probability on all prices within the original range of potential valuations $P \in [V_{\min}, V_{\max}]$. Increasing the accuracy of these beliefs therefore involves increasing $\bar{f}_\alpha(V_\beta)$, the ex ante probability placed on the true buyer valuation, while decreasing the probability placed on all other valuations. Similarly, sellers may differ in terms of their risk preferences, and changes to these preferences may result in sellers becoming either more risk seeking or more risk averse over the entire range of their respective function.

A large-enough increase in risk seeking preferences or the accuracy of beliefs will result in a permanent convention price that is at least as large as the resulting price without the increase. Holding the ex ante beliefs of all sellers and the utility functions of all but any one seller constant, it is possible to increase the risk-seeking preferences of the one seller such that the permanent convention price equals the buyer valuation V_β . Similarly, holding the utility functions of all sellers and the ex ante beliefs of all except any one seller constant, it is possible to increase the accuracy of the one seller's ex ante beliefs regarding the buyer valuation such that the permanent convention price equals the buyer valuation V_β . These two results hold regardless of which seller types are drawn to play the stage game in each period leading up to the establishment of the convention. Note also that depending on the original characteristics of the sellers, these results may require a very extreme change to the agent's ex ante beliefs or risk preferences. Relaxing the significance of the change in seller characteristic can still result in a guaranteed convention price at least as close to the buyer valuation as arises without the change, assuming that the change remains significant

enough.¹⁰

3.3 Limiting Observations to Recent Transactions

The previous section considered the model when sellers have perfect knowledge of the transaction histories. This section weakens this assumption, and considers a case where sellers only have knowledge of the most recent m periods. Because we are concerned with comparing the differences in agent interactions that result when class A agents have finite-period memory compared to full knowledge of the game, m is assumed to be *sufficiently large* to make such comparisons reasonable.¹¹

Similar to the analysis of the model under perfect knowledge of past transactions, this section considers how the process converges to conventions. However, there are some differences through which the process converges. When a player of type α is drawn to play the stage game in period t , she observes the most recent m transactions of the game. Therefore, the bounds on her potential price range are defined as the maximum *observed* accepted price (p_t^{\min}), and the minimum *observed* rejected price (p_t^{\max}); they are no longer determined by the entire history of the game, but just from those observed periods.

With perfect knowledge of past periods, sellers are aware of all past period price announcements, and therefore the potential price range can only converge, or remain constant over time. With limited memory (imperfect knowledge) of past period transactions, a price announcement is forgotten m periods after it is announced. When a bound on the period- t potential price range is not observed by the seller in the following period, $t + 1$ (meaning the bound must have been announced in period $t - m$), the $t + 1$ seller necessarily faces a different updated belief distribution than she would have if the bound was observed. This

¹⁰In general, increasing the accuracy of ex ante seller beliefs regarding buyer valuation or seller risk seeking can result in a permanent convention price that is less than the convention price achieved before the increase.

¹¹Any $m \geq \frac{P_{\max} - V_{\beta}}{\delta}$ is always sufficiently large for all claims in this paper to hold; however, m usually can be much smaller than that, depending on the specifics of the model parameters. Requiring a sufficiently large m assures that p^{\min} is not forgotten as the sellers experiment with other prices within a potential price range.

dynamic has the potential to significantly change the path of convergence over the course of the game.

In analyzing this issue, we first establish that the lower bound of the potential price range is never forgotten. This follows almost directly from the assumption that memory remains sufficiently large. Suppose that a price announcement is accepted in period t . Therefore, this price is the lower bound on the potential price range in period $t + 1$. Large-enough memory ensures that in some period $s \in [t + 1, t + m]$, before p_t is no longer observed, a seller must either announce a higher price that is accepted, or announce a price $p_s = p_t$. In either case, the lower bound is not forgotten until after period $s + m$, and the reasoning repeats itself. Therefore, the lower bound on the potential price range will only converge toward the buyer valuation, or remain constant between periods.

Alternatively, the upper bound will eventually be forgotten, at which time the upper bound on the potential price range reverts back to $V^{\max} + \delta$. As long as a rejected past price announcement is observed by the sellers, it will not be announced again since sellers recognize that it results in a payoff of zero for sure. If an upper bound is maintained for m sequential periods, it is forgotten in the following period. This happens whenever all price announcements over the m periods are accepted. When the process achieves a convention the upper bound will necessarily be in place for m sequential periods. However, achieving a convention is not required for the upper bound to be forgotten.

The seller in the period in which the upper bound on the potential price range is forgotten then update her expectations regarding buyer valuation according to Equation 1, in which p^{\max} now equals $V_{\max} + \delta$. Remember that this implies that sellers are naive in the sense that when they observe no rejected values, they do not infer that some rejected prices may have been experienced in the past, then forgotten. A further discussion of this assumption is provided in Footnote 6.

Because of the different potential price ranges that arise under imperfect and perfect information, the path of convergence may differ between the two cases, even when the same

agent types are drawn to play the game in each period. This follows because the optimal price given a potential price range may not also be the optimal price if the upper bound on the potential price range is forgotten.

$$\begin{aligned} \tilde{p}_w &\in \arg \max u_\alpha(p) [1 - F_\alpha(p - \delta | p_w^{\min}, p_w^{\max})] \\ &\Rightarrow \tilde{p}_w \in \arg \max u_\alpha(p) [1 - F_\alpha(p - \delta | p_w^{\min}, V_{\max} + \delta)] \end{aligned} \quad (7)$$

With imperfect information of past periods, the process will converge to and achieve a convention, provided that seller memory is sufficiently large. The process converges to a convention in a similar fashion to the permanent convention in the previous section; however, a permanent convention is generally no longer achieved. Alternatively, we are concerned with repeatedly-reoccurring conventions, as defined above.

When a convention is eventually established, limits to seller memory mean that the process will not generally remain in the convention indefinitely. When the upper bound of the potential price range is no longer observed, the process generally leaves the convention. The process will again converge to another convention, that may or may not be at the same price. However, eventually, the process will achieve a convention \hat{p} , such that after the process leaves the convention, it necessarily re-converges to a new convention at the same price \hat{p} . This necessary re-convergence will continue to be present the following time the process leaves the convention, and the process therefore enters a cycle in which the same convention price reoccurs on an ongoing basis throughout the remainder of the game.

Proposition 8 *There exists a value s such that for any memory length $m > s$, the process almost surely converges to a repeatedly-reoccurring convention, p_{rrc}^* .*

Intuition for this result is provided here. As discussed above, long-enough memory is assumed such that sellers can experiment with higher prices (which may be rejected) without forgetting the value of the most-recent accepted price. This means that the lower bound of the potential price range is never forgotten, and can therefore only converge or remain

constant. Given a lower bound of a potential price range p^{\min} , suppose that starting from a potential price range of $[p^{\min}, V_{\max}]$ there exists some draw order of sellers that results in the lower bound of the potential price range converging to a higher price. Although this convergence is not necessarily achieved by any point in time, with an infinite time horizon, if convergence is possible, the lower bound will eventually converge. Since the lower bound can only converge a finite number of times, the process eventually achieves a state in which further convergence of the lower bound is not possible. We label this point p_{rrc}^* . When this happens, no price higher than p_{rrc}^* will be accepted in any future period. The process will fluctuate between having p_{rrc}^* announced and accepted, and having different, higher prices announced and rejected. If a higher price is accepted, the lower bound will converge further, which is a contradiction; therefore, no higher price can be accepted. When these price experiments are rejected, the upper bound on the potential price range converges. If no experimentation take place for m periods, that means that the upper bound has been established for m period, and in the following period it reverts back to $V_{\max} + \delta$. Then, more experimentation can take place, and the upper bound again converges. This cycle repeats through the duration of the game.

As this cycle repeats indefinitely, p_{rrc}^* is the only price that is ever announced and accepted. Notice that a convention involving p_{rrc}^* may not generally be established prior to an established upper price range bound being forgotten and the potential price range reverting back to $[p_{rrc}^*, V^{\max}]$. However, a convention is achieved at least occasionally when the order of seller draws cause the upper bound to converge such that all sellers find it optimal to choose p_{rrc}^* . Given our assumptions, this is always possible, and, with an infinite time horizon, will happen on a reoccurring basis. It is important to recognize that the process does not have to be in a convention in order for sellers to announce p_{rrc}^* . This is because the drawn sellers may prefer to announce the price, even if it is not the optimal price for *all* members of the class of sellers. In contrast, a convention is only officially established when all sellers agree on the optimal price.

Once the process achieves a repeatedly-reoccurring convention, it follows that, as $t \rightarrow \infty$, the percent of time that p_{rrc}^* is announced approaches some constant, λ ; and the portion of time that a price is announced and rejected approaches $1 - \lambda$. Therefore, whether the repeatedly-reoccurring convention price can make the average seller better off depends on both the relative level of p_{rrc}^* compared to p_{pc}^* , and the parameter λ . This potential benefit is discussed in more detail in another section below.

The position of the repeatedly-reoccurring convention price p_{rrc}^* relative to V_β and p_{pc}^* depends on the order that sellers are randomly drawn to play the period games. The repeatedly-reoccurring convention price may be closer to the buyer valuation than the permanent convention price that would have been established if all agents had perfect information regarding past transactions. The following two propositions provide sufficient conditions for p_{rrc}^* to be at least as large as p_{pc}^* ; however, these are not necessary conditions, and p_{rrc}^* greater than p_{pc}^* can still result even if these conditions are violated.

Proposition 9 *Given the order of the seller draws, there exists a value w such that for any memory length $m > w$, the process almost surely converges to a repeatedly-reoccurring convention p_{rrc}^* , such that $p_{pc}^* \leq p_{rrc}^* \leq V_\beta$, where p_{pc}^* is the permanent convention price that would have been established with perfect information.¹²*

If we are certain regarding the seller type that will be drawn to play the game in each period, we can determine the period in which a process will first achieve a permanent convention if all sellers had perfect information of past transactions. Denote this period by \tilde{w} .

¹²This proposition always holds ex post, in that once we know the period in which a permanent convention is established, we can determine the minimum value of w for the proposition to hold. However, because sellers are drawn randomly in each period, we can not generally be certain of the order of seller draws, and can not determine the value of w prior to a game being played. The exception to this is when the class of sellers is homogenous, or, equivalently, when the class is composed of a single agent. In this case, we are certain as to the seller type that plays the game in any given period, and can determine the minimum value of w prior to the game. With a heterogeneous seller class, it is possible to draw the same seller type for any finite number of periods in a row. For any w that we set prior to the start of the game, there is a positive probability that the upper bound of the potential price range is forgotten prior to the lower bound of the potential price range converges to p_{pc}^* . If this happens, it is possible for the process to achieve a repeatedly-reoccurring convention price less than p_{pc}^* . However, if the random draw of sellers is not given, other conditions can help ensure that the repeatedly-reoccurring convention is at least as close to V_β as the alternative permanent convention.

Any memory length greater than \tilde{w} ensures that the lower bound on the potential price range converges to p_{pc}^* before the consequences of limited memory come into play. Additionally, as we previously discussed, setting memory greater than $(V_{\max} - V_{\beta}) \frac{1}{\delta}$ is always sufficient to assure that the lower bound of the potential price range can only converge to a higher value or remain constant. Therefore, it is possible to pick a memory length w high enough to meet both these requirements. For such a w , once the lower bound achieves the value p_{pc}^* , any future accepted price, including p_{rrc}^* , will be at least as large as p_{pc}^* . After the establishment of p_{pc}^* , limiting memory will cause the upper bound to occasionally be forgotten, which causes the potential price range to expand to $[p_{pc}^*, V_{\max}]$. When this happens, sellers are faced with additional potential prices. Depending on their risk preferences and ex ante expectations regarding buyer valuation, they may decide to experiment further with price. If any of these further price experimentations are excepted, then the lower bound on the potential price range converges further, and any future convention will be at a strictly-higher price than p_{pc}^* .

Even when sellers are not identical, the sellers may be similar enough in terms of their ex ante beliefs and risk preferences that, under perfect memory, the process always achieves a permanent convention by some fixed period. Such a seller class is called *quasi-homogeneous*.

Definition 10 *A class of sellers is **quasi-homogeneous** if under perfect memory, and the random draw of sellers, there exists some finite \hat{T} , such that for any $m \geq \hat{T}$, the process achieves a permanent convention in the first s periods of the game with probability one.*

When the class of sellers is quasi-homogeneous, it is possible to fix a memory length large enough such that the lower bound on the potential price range will always converge to p_{pc}^* before the upper bound is forgotten.

Proposition 11 *If the class of sellers is quasi-homogeneous, there exists a value s such that for any memory length $m > s$, the process almost surely converges to a repeatedly-reoccurring convention p_{rrc}^* , such that $p_{pc}^* \leq p_{rrc}^* \leq V_{\beta}$, where p_{pc}^* is the permanent convention price that would have been established with perfect information.*

The more similar the sellers are in terms of risk preferences and expectations, the more likely a class of sellers is quasi-homogeneous. Whenever the class of sellers is homogeneous, this condition always holds, and, with long-enough memory, the repeatedly-reoccurring convention always results in a price at least as close to the buyer valuation as the alternative permanent convention that would result under perfect information.

As these similar propositions show, limiting seller information to the m most-recent periods can achieve a long-term convention price at least as close to the buyer valuation as under perfect information of game history. Because of the random draw of sellers from a heterogeneous seller class, determining the minimum memory length to achieve this result is generally not feasible, and the result only holds with certainty in the limit. As memory increases, but remains finite, the probability that the result holds increases. However, even though the results only hold with certainty in the limit, these results hold with positive probability for any m sufficiently large. For any memory length above this limit, limiting observations to recent periods has the potential to result in at least as large of a long-term convention price, and the potential to increase long-term average seller and class revenues.

3.4 Random Observations of Past Periods

In this section, we no longer limit seller observations to the most recent periods, and instead allow sellers to observe each of the past transactions with some independent positive probability $\rho \in (0, 1)$. ρ may be constant across all past periods, or it may be decreasing with distance in time from the current period, such that the more recent the past transaction, the more likely a seller is to observe it. In each period, there is a positive probability of observing no past transactions, all of the past transactions, or any incomplete selection of past periods. Therefore, the process does not generally converge as it did under consecutive period observations.

Sellers are not aware of the expected values, or the total number of past observations; they are only aware of the observations they do observe. There always exists a positive

probability that any price that was previously announced, including p_{t-1}^{\min} or p_{t-1}^{\max} , will not be observed in period t . There is no assurance the potential price range will either remain constant, or converge. Additionally, if the process converges between periods $t - 1$ and t , there is no longer the requirement that either $p_t^{\min} = p_{t-1}$ or $p_t^{\max} = p_{t-1}$. Therefore, instead of achieving a permanent or repeatedly-reoccurring convention after some finite number of periods, the process achieves a *convention in probability* as $t \rightarrow \infty$.

When the process does achieve a convention in probability at price p_{cip}^* , then as $t \rightarrow \infty$, the probability that $p_t^{\min} = p_{cip}^*$ approaches 1 in the limit, and the probability that $p_t^{\max} \in \{p^{\max} \mid p_{cip}^* \in \arg \max_p u_\alpha(p) [1 - F_{\alpha t}(p - \delta \mid p_{cip}^*, p^{\max})]\}$ for all $\alpha \in A$ also approaches 1 in the limit. Therefore, in the long run, the convention price p_{cip}^* is announced almost all of the time.

Proposition 12 *When sellers observe each past period with constant probability ρ , the process almost surely achieves a convention in probability at price p_{cip}^* as $t \rightarrow \infty$.*

The formal proof is reserved for the appendix; however, intuition is provided here. Note that the more often a price was announced in periods 1 through $t - 1$, the higher the probability that it is observed by the seller at time t . The analysis depends on the likelihood that a price becomes p_t^{\min} or p_t^{\max} in any period t . A past accepted price becomes p_t^{\min} when it is observed by the seller in period t , and no higher past accepted price is also observed. The probability that any past price is observed is increasing in the number of past periods in which it was announced. Therefore, when a price is accepted in period t it reduces the probability that any lower price will be seen as p_{t+1}^{\min} . Only the largest accepted price announced over the course of the game (as $t \rightarrow \infty$) is not subject to its probability of being p_t^{\min} being reduced. We denote this largest accepted price by p_{cip}^* . It follows that as $t \rightarrow \infty$, the probability that $p_t^{\min} = p_{cip}^*$ goes to 1.

The upper bound on the potential price range evolves in the limit similar to p_t^{\min} . The probability that a past rejected price is p_t^{\max} depends on the probability that the price is observed, and the probability that any lower previously rejected price is also observed.

Therefore, when a price is rejected in period t it reduces the probability that a higher price will be p_{t+1}^{\max} . Only those past rejected prices that when observed cause all sellers $\alpha \in A$ to announce a price below the buyer valuation are not necessarily subject to their probabilities of being p_t^{\max} being reduced as $t \rightarrow \infty$. In the long run, the probability that p_t^{\max} equals one of these immune past rejected prices goes to 1.

It is easy to see that $p_{cip}^* \geq p_{pc}^*$ is possible. Because it is possible to observe any subset of all past transactions, it is feasible that the seller observe the highest past accepted price and the lowest past rejected price during the first τ periods of the game. As τ increases, this becomes less likely, but is still possible. In this situation, the process converges just as it would under perfect information during the first τ periods. For τ large enough, the process will establish price p_{pc}^* prior to period τ . Since no long term convention can be established at a price less than the highest past accepted price, it follows that p_{cip}^* will be at least as large as p_{pc}^* in this case. Therefore, $p_{cip}^* \geq p_{pc}^*$ is possible. If observing some other p^{\min} and p^{\max} in a period following τ causes sellers to experiment with price, they may eventually announce a new price between p_{pc}^* and the buyer valuation. If that new price is accepted, then a p_{cip}^* strictly greater than p_{pc}^* is assured.

Proposition 13 *There is a strictly positive probability that the convention in probability price established for any $\rho \in (0, 1)$ is at least as large as the permanent convention price established when $\rho = 1$. Formally, $\Pr \{p_{cip}^* \geq p_{pc}^*\} > 0$.*

Without adopting further assumptions regarding the structure of the sellers' utility functions (including risk preferences), pre-game beliefs regarding the buyer valuations, or the size of the original potential price range, we cannot draw more exact conclusions about the likelihood the process settles to a convention in probability price greater than the permanent convention price that would have been established under perfect information.¹³ Generally, the convention in probability price that a process may settle to ex ante is not unique. De-

¹³Further research may impose some further structure on the game, and simulate the long term results.

spite this, we can determine some characteristics of the set of possible ex ante convention in probability prices.

Let C denote the set of conventions that for all $p^* \in C$:

1. p^* may be achieved following the initial game period in which the price history is an empty set;
2. p^* may be the maximum past rejected price when there first exists some p^{\max} within the set of past rejected prices, such that $[p^*, p^{\max})$ would form a convention under perfect information; and
3. for all p^{\max} within $H^R(p^*)$, a period- t potential price range $[p^*, p^{\max})$ necessarily implies that either $p_t = p^*$ or $p_t \in H^R(p^*)$. $H^R(p^*)$ is the set of rejected past prices as $t \rightarrow \infty$ when the convention in probability settles to p^* .

Let $\Omega(p^*)$ denote the ex ante probability that the process achieves a convention in probability price equal to p^* .¹⁴ Ex ante, there is a positive probability that any $p^* \in C$ may be established as the convention in probability price. In fact, $\Omega(p^*) > 0$ if and only if $p^* \in C$. Also, p_{pc}^* is contained in C for sure. Long-run expected average seller utility is therefore given by

$$\sum_{\alpha \in A} \pi(\alpha) \sum_{p^* \in C} \Omega(p^*) u_{\alpha}(p^*) \quad (8)$$

¹⁴Consider a graph depicting the potential price path as established by the Markov process defined in the model as the dynamic price process. Each node in the graph represents the price history achieved up until that node is reached. The nodes are connected by edges, representing the potential evolution of the complete price history between periods. Nodes may be classified by the period of play in which they occur, and each period- t node can only have one predecessor node in each of earlier periods $1, \dots, t-1$. Therefore, there is a unique path from any node back to the initial node of the game. There may be multiple nodes in representing the same price history, since the same price history may follow from different past histories of play. The weight of the edges connecting the nodes represents the probability that a node follows from its immediate predecessor. The probability that a node is observed during the play of a game is the product of the weights of the edges leading from the node along the unique path of play back to the initial period. The probability that any game history is observed is the sum of these probabilities across all nodes representing the game history. In the long run, as established by the previous convention, the process approaches a convention in probability. When $\rho \in (0, 1)$, the graph is infinite. However, if we consider the graph through a large t , the majority of paths of play within the graph will clearly approach conventions (where one accepted price, and one rejected price are announced often enough such that they are observed almost all the time). If we know the functional forms of the sellers utility functions, and belief preferences, as well as the valuation, we can estimate Ω for each $p^* \in C$.

Alternatively, ρ may differ across past periods, such that sellers are more likely to observe recent periods, compared to periods long past. In this case, $\rho(\Upsilon)$ is a continuous function assigning a probability to each of the past periods, where the probability represents to likelihood that the seller observes the past period. Υ is the number of period difference between the current period and the past period for which the function assigns the probability. If the current period is t , then $\Upsilon = 1$ represents period $t - 1$, and $\Upsilon = n$ represents period $t - n$. $\rho'(\Upsilon) > 0$ for all $\Upsilon = 1, 2, 3, \dots$, $\rho(1) \in (0, 1)$, and $\rho(\Upsilon) \rightarrow \xi$ as $\Upsilon \rightarrow \infty$, where $\xi \in [0, 1)$. When $\xi = 0$, the probability that a seller observes a past period approaches zero as that period becomes farther from the present period. Alternatively, increases in time passed may cause the probability of observing a past period to approach some positive number when $\xi \in (0, 1)$. It can be shown that when sellers observe past periods according to function $\rho(\Upsilon)$, and when $\xi \in (0, 1)$, the process almost surely converges to a convention in probability as $t \rightarrow \infty$.

If, alternatively, $\xi = 0$, then the process continues to achieve the convention in probability price \hat{p}^{\min} part of the time. However, with decreasing probability associated with observing the upper potential price range bound, there will always be some level of experimentation, even in the limit, as the probability of observing \hat{p}^{\max} decreases with time, until it is announced again. In this situation, the welfare benefits of limited information may continue to hold, so long as $\hat{p}^{\min} > p_{pc}^{\min}$ and the portion of time that the process achieves price \hat{p}^{\min} is high enough to out weigh the potential welfare loss during periods of experimentation.

3.5 Long-Term Benefits of Imperfect Information

As shown above, both types of memory limits can result in long-run prices closer to the buyer valuation than would result if all sellers had perfect knowledge of all past transactions. These higher long-run prices can mean improved long-term average seller welfare. Long-run average

seller utility is given generally by

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t u_{\alpha_t}(p_t) a_t \quad (9)$$

Under perfect information of game history, this becomes

$$\sum_{\alpha \in A} \pi(\alpha) u_{\alpha}(p_{pc}^*) \quad (10)$$

Where p_{pc}^* is the permanent convention price, and $\pi(\alpha)$ is the probability that seller α is drawn to play the stage game in any period.

Alternatively, if memory is limited to the most recent m periods, average seller utility becomes

$$\sum_{\alpha \in A} \lambda_{\alpha} \pi(\alpha) u_{\alpha}(p_{rrc}^*) \quad (11)$$

Where p_{rrc}^* is the price associated with the repeatedly-reoccurring convention, and λ_{α} represents the proportion of time, as $t \rightarrow \infty$, that $\alpha \in A$ announces p_{rrc}^* when she is selected to play the game. λ_{α} therefore represents the proportion of time the process is in the repeatedly-reoccurring convention, plus the proportion of time the process is not in the repeatedly-reoccurring convention but agent α finds it optimal to announce p_{rrc}^* regardless. In the limit, as $t \rightarrow \infty$, λ_{α} approaches a constant on $(0, 1]$ for all $\alpha \in A$. Since the process necessarily leaves the convention regularly over the course of the game, by definition of a repeatedly-reoccurring convention, there must exist at least one $\alpha \in A$ such that $\lambda_{\alpha} \neq 1$. λ without the subscript denotes the total portion of time, as $t \rightarrow \infty$, that p_{rrc}^* is announced by any agent. Similarly, this value will approach a constant.

If each past transaction has a probability $\rho \in (0, 1)$ of being observed in any period, the average seller utility function becomes

$$\sum_{\alpha \in A} \pi(\alpha) u_{\alpha}(p_{cip}^*) \quad (12)$$

Where p_{cip}^* denotes the price associated with the convention in probability.

Average long-term seller utility may be higher when memory is limited to the most recent m periods of the game, compared to when sellers have perfect knowledge of all past periods iff

$$\sum_{\alpha \in A} \lambda_{\alpha} \pi(\alpha) u_{\alpha}(p_{rrc}^*) > \sum_{\alpha \in A} \pi(\alpha) u_{\alpha}(p_{pc}^*) \quad (13)$$

Which is clearly possible, but not necessarily the case. As $m \rightarrow \infty$, but remains finite, $\lambda_{\alpha} \rightarrow 1$ for all $\alpha \in A$, and $\Pr(p_{rrc}^* \geq p_{pc}^*) \rightarrow 1$. Therefore, this requirement will hold at least with equality in the limit. If it holds outside of the limit depends not only on the model parameters, but also on the random draw of agents.

When $p_{cip}^* \geq p_{pc}^*$, it is always true that

$$\sum_{\alpha \in A} \pi(\alpha) u_{\alpha}(p_{cip}^*) \geq \sum_{\alpha \in A} \pi(\alpha) u_{\alpha}(p_{pc}^*) \quad (14)$$

When $p_{cip}^* > p_{pc}^*$, in the long run, the average seller is strictly better off with under this type of limited memory.

These results show that under reasonable conditions, at the onset of the game, sellers may prefer limited, rather than perfect, information of past transactions to be shared with the randomly selected seller in each stage. Although limited information may result in some sellers announcing prices that had been rejected by buyers in the past and therefore receiving zero payoff in that period, it may also result in the majority of future seller prices being closer to the buyers' valuation. On the other hand, buyers may prefer the sellers to have perfect information regarding the history of play since it may result in a lower average sale price in the long-run. This suggests a potential role for consumer advocacy groups to assure that information regarding past transactions remains available when the sellers themselves do not find it optimal to implement a system of perfect information sharing.

3.6 Class of Sellers as a Firm

In addition to providing insight regarding the decision process of the individual sellers, the model can also provide insight regarding firm and employee relationships. Consider class A as a whole to represent a firm, and the individual members of the class as employees, sales people, or managers within the firm.¹⁵

Where sellers are assumed to only care about the current period, we allow the firms to place as much weight on future periods as it does on the present period. Firms are risk neutral in revenue, therefore firm welfare is given by average sales revenue over all periods of the game:

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t p_n a_n \quad (15)$$

We do not directly model the strategic interaction between the firm and the sellers; however, we acknowledge that such a relationship exists. Firms can influence employee risk preferences through the design of promotion and compensation agreements (see for example Wilson 1968, and Ross 1973). Additionally, firms may be able to influence seller risk preferences on a temporary basis through the use of sales contests, quotas, or promotion rules (Gaba and Kalra 1999). Considering such results in the context of our model, with perfect knowledge of game history, implementing a contest or other scheme to temporarily increase seller risk taking will never result in a convention price lower than the original convention price, so long as the scheme is implemented after the process has already converged to a potential price range close to the permanent convention price. Because of this, a firm that is able to temporarily increase employee risk taking may be able to assure long-term revenue at least as high as it would have been without this ability. Therefore, our model supports the selective use of sales contests and other means of increasing employee risk taking by the firm, and shows that such tools can increase long-term firm revenues. This results helps

¹⁵Alternatively, class A could represent any parent organization who's overall wellbeing is dependent upon the actions of its individual members. For example, class A may be a trade organization, and each $\alpha \in A$ may represent a firm with membership in that trade organization; or class A may represent a family or club, and each $\alpha \in A$ could be family or club members.

justify behavior that is frequently observed of actual firms.

Our model also allows us to consider the impact that a firm may have on long-term revenue if it is able to manipulate price choices by limited seller access to information regarding past transactions. We assume that the firm can only commit to long-term information sharing policies, and cannot choose how much information to provide on an individual basis.¹⁶ Firms can therefore choose policies of information sharing that remain in place for the duration of the game. In the previous sections of this paper, we consider how the process converges under perfect information of game history, and compare it with cases in which sellers are only aware of transactions in the most recent m periods, and where sellers randomly observe each past transaction. Here we assume that the firm can commit to policies that provide any one of these three types of information access.

Under perfect information, the firm's long-term welfare function is given by

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t p_{pc}^* = p_{pc}^* \quad (16)$$

Alternatively, under m -period memory, the firm's long-term welfare function is given by

$$\lim_{t \rightarrow \infty} \frac{\lambda}{t} \sum_{n=1}^t p_{rrc}^* = \lambda p_{rrc}^* \quad (17)$$

Under random past period observations, the firm's long-term welfare function is given by

$$\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t p_{cip}^* = p_{cip}^* \quad (18)$$

As already shown, $p_{rrc}^* \geq p_{pc}^*$ and $p_{cip}^* \geq p_{pc}^*$ are possible. Therefore, the firm's long-term welfare can be higher under either of the types of limited memory. For large enough m , the

¹⁶This assumption is justified given that the firm represents stock holders who do not typically have the ability to micro-manage firm operations. If firms are able to choose information on a per-transaction basis, then they would be able to even better use information limits to achieve higher long-run profits. This is because firms could manipulate the process similar to the case when sellers have random access to past period transactions, but firms could also avoid the periods in which prices lower than previously accepted prices are unnecessarily announced.

firm is never worse off under m -period memory compared to perfect seller information; and the firm expects to be at least as well off under random past period memory compared to perfect seller information when

$$Ep_{cip}^* = \sum_{p^* \in C} \Omega(p^*) p^* \geq p_{pc}^* \quad (19)$$

Therefore, employee information limits provide an alternative to more standard compensation tools such as sales contests, quotas, or promotion rules for firms to influence price experimentation amongst their employees. Where the more standard tools have direct costs associated with them, limiting information is costly in terms of potential lost sales as sellers experiment with prices greater than the buyer valuation.¹⁷ In the long run, with an infinite time horizon, the lost-sales cost is virtually eliminated in the case of random memory. However, when firms do not actually face an infinite time horizon, considering these costs is important. Even with these costs, limiting information may be the less costly means of increasing the long-term price compared to the use of more standard compensation tools.

3.7 Multi-Period Sellers

Up to this point, the paper has assumed that individual sellers ignore the effect their price announcements have on future-period payoffs. This may result if sellers only play the game once and are replaced in the set A by identical agents after play, or if sellers completely discount future period utility (discount rate $\eta = 0$). However, the results presented in this paper may continue to hold when sellers play the game more than once and do not completely discount future-period utility.

Let N denote the number of agents within class A , and $\eta_\alpha \in [0, 1]$ denote the discount rate an agent of type α applies to future period utility. $\eta_\alpha = 0$ implies that agent α completely

¹⁷Direct costs may include prizes for contests, hiring and training costs associated with replacing those who do not meet high quotas or promotion requirements, and others. Lost sales costs are likely to occur when standard compensation tools increase seller risk taking, as well as when memory is limited.

ignores future periods. When a seller of type α is selected to play in period t , the agent solves:

$$\max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p_t^{\min}, p_t^{\max})] + B(\eta_\alpha, N, p) \quad (20)$$

where $B(\eta, N, p)$ represents the expected, discounted increase in utility in all future periods of the game from announcing price p instead of the price p° , where p° is the price that α announces when future period payoff is ignored. Technically, $p^\circ \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p_t^{\min}, p_t^{\max})]$. Under the structure of the game previously described, it follows that $B(\eta, N, p^\circ) = 0$. More generally, $B(\eta, N, p) > 0$ if and only if $\eta \in (0, 1]$ and price p results in higher expected future period payoff compared to p° . Additionally, $B(0, N, p) = 0$; $\frac{\partial B}{\partial \eta} > 0$; $B(\eta, N, p) \rightarrow 0$ as $N \rightarrow \infty$; and $\frac{\partial B}{\partial N} < 0$.¹⁸

As $\eta \rightarrow 0$ or $N \rightarrow \infty$, all of the results in previously established continue to hold. However, considering these values in their limit is not required to maintain the results. So long as N is sufficiently large, or η is sufficiently small, the results continue to hold.

When N is finite, and $\eta \in (0, 1]$, announcing price $p \in (p^{\min}, p^{\max})$ will always result in at least as high of expected future utility compared with announcing price $p = p^{\min}$. This is because choosing $p \in (p^{\min}, p^{\max})$ can result in a smaller potential price range in future periods, decreasing uncertainty regarding buyer valuation, and potentially causing the process to converge to long-term convention price $\tilde{p} > p^*$, where p^* is the long-term convention price associated with the original model assumptions described above. However, this does not imply that $\tilde{p} = V_\beta$, only that $\tilde{p} \in [p^*, V_\beta]$. When $\tilde{p} < V_\beta$, the sellers can still receive additional long-term payoff improvements from increasing class risk aversion or introducing imperfect information.

With other factors held constant, an increase in N or a decrease in η results in an increase in the range of possible functional forms of $\{\bar{F}_\alpha(\cdot), u_\alpha(\cdot)\}_{\alpha \in A}$ such that limiting seller information may result in a higher long-term average sales price. It follows that

¹⁸This assumes that the probability that any individual agent is select to play in any given period is strictly decreasing in N .

allowing for a finite set of sellers who care about the strategic consequences of their price announcements does not change the model's fundamental results.

4 Concluding Remarks

Applying learning dynamics to a simple price-setting game provides a framework to analyze seller-buyer market transactions, where sellers learn about willingness to pay over time. Our analysis shows how the long-term price depends on the sellers' risk preferences and ex ante beliefs regarding the buyer valuation, which supports the use of sales contests, quotas, and promotion schemes by firms in an attempt to temporarily increase risk taking amongst its employees. Additionally, the long-term transaction price depends on seller memory, or access to information regarding past transactions. Surprisingly, limiting seller memory can result in higher long-term prices, and increase average seller utility. Sellers and firms may therefore exhibit ex ante preferences for a game in which sellers have imperfect rather than perfect information of game history.

In addition to the questions addressed in this paper, our model may be expanded to explore additional issues regarding buyer-seller interactions. There are opportunities to expand the model to include more complex strategies. One such expansion involves developing a more detailed analysis of firm-seller interaction in the framework of this model. Such an expansion would allow for a greater consideration of firm decisions to implement sales contests, quotas, or certain promotion rules. Another expansion may allow for strategic buyers, which would result in a more complex relationship between the sellers and the buyers. Furthermore, our model currently assumes that the product or service provides the same benefit to all buyers, and that memory is *large enough*. It would be interesting to consider how the process behaves if the class of buyers is heterogeneous or seller memory is very short; specifically, what conditions allow for the conclusions in this paper to be generalized for these alternative cases?

References

- [1] Basu, A.K., R. Lal, V. Srinivasan, and R. Staelin, 1985, "Salesforce compensation plans: an agency theoretic perspective," *Marketing Science*, 4n4, 267-291.
- [2] Gaba A., and A. Kalra, 1999, "Risk behavior in response to quotas and contests," *Marketing Science*, 18n3, 417-435.
- [3] Holmstrom, B. and M.P. Page, 1999, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66n1, 169-182.
- [4] Hurkens, S., 1995, "Learning by Forgetful Players," *Games and Economic Behavior*, 11, 304-329.
- [5] Ross, S.A., 1973, "The Economic Theory of Agency: The Principal's Problem," *American Economic Review*, 63n2, 134-139.
- [6] Weibull, J., 1995, *Evolutionary Game Theory*, The MIT Press, Cambridge, MA.
- [7] Wilson, R., 1968, "The Theory of Syndicates," *Econometrica*, 36n1, 119-132,
- [8] Young, H.P., 1993a, "The Evolution of Conventions," *Econometrica*, 61n1, 57-84.
- [9] Young, H.P., 1993b, "An Evolutionary Model of Bargaining," *Journal of Economic Theory*, 59, 145-168.
- [10] Young, H.P., 2005, *Strategic Learning And Its Limits: Arne Ryde Memorial Lectures 2002*, Oxford University Press, UK.

5 Proofs

Proposition 14 Proof of proposition 6. *Perfect information implies that the potential price range can only converge or remain constant between periods. If at time t , an $\alpha \in A$ is drawn such that $p_t^{\min} \notin \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p_t^{\min}, p_t^{\max})]$, then the process must converge. If an alternative seller $\hat{\alpha}$ is drawn, such that $p_t^{\min} \in \arg \max_p u_{\hat{\alpha}}(p) [1 - F_{\hat{\alpha}}(p - \delta | p_t^{\min}, p_t^{\max})]$,*

the process does not converge and the same information is available to the next seller in period $t + 1$. Suppose the process is not in a convention at time t , then there does exist at least one α such that $p_t^{\min} \notin \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p_t^{\min}, p_t^{\max})]$. As $t \rightarrow \infty$, such an α will eventually be drawn, and the process will converge. However, the process can only converge at most $(V_{\max} - V_{\min}) \frac{1}{\delta}$ times before either $p^{\min} = p^{\max} - \delta$ or otherwise there does not exist an $\alpha \in A$ such that $p^{\min} \notin \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | p^{\min}, p^{\max})]$. In either case, the process achieves a convention. Since the process only converges, or remains constant, it remains there indefinitely. ■

Proof of proposition 7. Let the potential price range be $[V_{\min}, V_{\max}]$, and fix any seller risk preferences. For all seller types $\alpha \in A$, increase $f_\alpha(V_{\min})$ while simultaneously decreasing the values $f_\alpha(P)$ for all $P \in (V_{\min}, V_{\max}]$. As we do this, $f_\alpha(V_{\min}) \rightarrow 1$, and for all $P \in (V_{\min}, V_{\max}]$, $f_\alpha(P) \rightarrow 0$. Continue these transformations keeping $f_\alpha(V_{\min}) < 1$ and $f_\alpha(P) > 0$ for all α . Since risk preferences are fixed, sellers eventually become sure enough that $V_\beta = V_{\min}$ that even the most risk seeking seller selects $p = V_{\min}$. When this happens, $V_{\min} \in \arg \max_p u_\alpha(p) [1 - F_\alpha(p - \delta | V_{\min}, V_{\max} + \delta)]$ for all $\alpha \in A$, and a permanent convention is established at price $V_{\min} < V_\beta$.¹⁹ ■

Proof of proposition 8. Let $m \geq (V_{\max} - V_\beta) \frac{1}{\delta}$, which ensures that the lower bound of the potential price range is never forgotten.

Note that the upper bound of the potential price range will eventually be forgotten. A rejected price announcement, so long as it is remembered, will not be played again. Therefore, a rejected price announcement is forgotten $m + 1$ periods after it is announced. If a upper bound on the potential price range remains in place for m periods, in the following period, it is forgotten and the upper bound becomes $V_{\max} + \delta$. For the potential price range upper bound to not become $V_{\max} + \delta$, after an upper bound is established, a new one must replace it within m periods. However, there can exist at most $(V_{\max} - V_\beta) \frac{1}{\delta} - 1$ different upper bound prices. It follows that the process can only last at most $m [(V_{\max} - V_\beta) \frac{1}{\delta} - 1]$ periods before

¹⁹Alternatively, this proposition can be changed to fix the agents' ex ante beliefs, and then increase the risk aversion of sellers until a similar result is established.

the upper bound is forgotten, and becomes $V_{\max} + \delta$.

Let $\Xi(p^{\min})$ define a graph representing the potential price ranges that may follow from an initial potential price range $[p^{\min}, V_{\max}]$. Each vertex, or node, in the graph represents a potential price range that may be established. Edges connect the initial node, representing $[p^{\min}, V_{\max}]$, to each potential price range that could potentially be established in the period immediately following $[p^{\min}, V_{\max}]$, with the exception of $[p^{\min}, V_{\max}]$ itself. These new vertices are then directly connected by an edge to new vertices representing each of the potential price ranges that may immediately follow them, with the exception of themselves and new ranges that are created by the reversion of the upper bound back to V_{\max} once it has already converged to a lower value. Vertices are only allowed to map back to one parent node; therefore, since potential price ranges may result following multiple parent nodes, there can be multiple vertices representing a single potential price range. When a potential price range will never result in further convergence to a new potential price range, it is represented by a terminal node of the graph, since there are no edges connecting it to sub-nodes. In other words, a terminal node represents a convention.

Let $\Gamma(p^{\min})$ represent the set of lower potential price range bounds included in the terminal nodes of the graph $\Xi(p^{\min})$, with the exclusion of p^{\min} itself. If $\Gamma(p^{\min}) \neq \emptyset$, then as $t \rightarrow \infty$, the process will eventually achieve convergence of the lower bound (even if the upper bound is forgotten multiple times before convergence eventually takes place). However, it is only possible for the lower bound to converge at most $(V_{\beta} - V_{\min}) \frac{1}{\delta}$ times over the course of the game. Therefore, as $t \rightarrow \infty$, the process will eventually achieve a lower bound such that $\Gamma(p^{\min}) = \emptyset$. When this happens, p^{\min} will remain the lower bound of the potential price range in all future periods. Although each convergence sequence does not necessarily result in a convention being established, it will happen whenever the process reaches a terminal node in $\Xi(p^{\min})$. The price associated with these conventions must be p^{\min} . ■

Proof of proposition 9. For any fixed order of seller draws, it is possible to determine with certainty the period in which the process would achieve a permanent convention under

perfect information of past transactions. Denote this period by \hat{T} . Set memory m such that $m \geq \hat{T}$ and $m \geq (V_{\max} - V_{\beta}) \frac{1}{\delta}$. Setting $m \geq \hat{T}$ implies that memory is sufficiently large such that the alternative permanent convention is established before periods are forgotten. Additionally, setting $m \geq (V_{\max} - V_{\beta}) \frac{1}{\delta}$ implies that the lower bound on a potential price range is never forgotten, and that the lower bound can only converge or remain constant between periods. Therefore, any convention established in periods following the alternative permanent convention must be at a price at least as large as the alternative permanent convention price.

■

Proof of proposition 11. Similar to previous proof. By definition of a quasi-homogenous class of sellers, there exists a memory length \hat{T} such that for $m \geq \hat{T}$, the process achieves a permanent convention before any periods are forgotten. Additionally, setting $m \geq (V_{\max} - V_{\beta}) \frac{1}{\delta}$ implies that the lower bound on a potential price range is never forgotten, and that the lower bound can only converge or remain constant between periods. Therefore, any convention established in periods following the alternative permanent convention must be at a price at least as large as the alternative permanent convention price. ■

Proof of proposition 12. As $t \rightarrow \infty$, let H^A denote the set of all past accepted prices as well as V^{\min} , and H^R denote the set of all past rejected prices as well as $(V^{\max} + \delta)$. Let $H = H^A \cup H^R$, noting that these sets only contain prices, and not formal price histories such as those that make up H_t . Additionally, define $\hat{H}^R = \{\hat{p} \mid p^{\min} \in \arg \max_p u_{\alpha}(p) [1 - F_{\alpha}(p - \delta \mid p^{\min}, \hat{p})]\}$ for all $\alpha \in A$ and $p^{\min} = \max \{p \in H^A\}$. Note that $\hat{p} \in \hat{H}^R$ are necessarily the lowest values in H^R .

For any period t , α_t chooses $p_t \in \arg \max_p u_{\alpha}(p) [1 - F_{\alpha}(p - \delta \mid p_t^{\min}, p_t^{\max})]$. If p_t is accepted, then the expectation regarding p_{t+1}^{\min} changes relative to the expectation of p_t^{\min} such that $\Pr [p_{t+1}^{\min} < p_t]$ decreases, $\Pr [p_{t+1}^{\min} = p_t]$ increases, and $\Pr [p_{t+1}^{\min} > p_t]$ remains unchanged; and the expectation regarding p_{t+1}^{\max} remains unchanged. If p_t is rejected, then the expectation regarding p_{t+1}^{\min} remains unchanged; and the expectation regarding p_{t+1}^{\max} changes relative to the expectation of p_t^{\max} such that $\Pr [p_{t+1}^{\max} > p_t]$ decreases, $\Pr [p_{t+1}^{\max} = p_t]$ increases,

and $\Pr [p_{t+1}^{\max} < p_t]$ remains unchanged. It follows that for all t , $E(p_{t+1}^{\min}) \geq E(p_t^{\min})$ and $E(p_{t+1}^{\max}) \leq E(p_t^{\max})$; and for $p_t \neq V_{\min}$, either $E(p_{t+1}^{\min}) > E(p_t^{\min})$ or $E(p_{t+1}^{\max}) < E(p_t^{\max})$.

After $p^{\min} = \max \{H^A\}$ is first announced as a price, it is observed by the seller in the following period with probability $\rho > 0$. Therefore, eventually, it is observed by a seller in some period, s . When it is observed, by construction, it must be p_s^{\min} . However, p_s^{\max} may either be in the set \hat{H}^R or $H^R \setminus \hat{H}^R$. If $p_s^{\max} \in \hat{H}^R$, then α_s selects $p_s = p^{\min}$, which increases the probability that p^{\min} is observed in any given future period (when p^{\min} was selected once previously, it is observed with probability ρ ; when it was selected in two previous periods, it is observed with probability $\rho^2 + 2(1 - \rho)\rho = 2\rho - \rho^2$). Furthermore, $p_s^{\max} \in \hat{H}^R$ means that $E(p_{s+1}^{\max}) = E(p_s^{\max})$. If $p_s^{\max} \in H^R \setminus \hat{H}^R$, then $p_s \in (p^{\min}, p_s^{\max})$, which means that in the following period $E(p_{s+1}^{\max}) < E(p_s^{\max})$; specifically, $\Pr [p_{t+1}^{\max} > p_t]$ decreases while $\Pr [p_{t+1}^{\max} = p_t]$ increases relative to the expectations in period t . As $t \rightarrow \infty$, this means that $\Pr [p_t^{\max} \in \hat{H}^R] \rightarrow 1$. Now, for large enough t , when $p_t^{\min} = \max \{p \in H^A\}$ is drawn, it is paired with $p_t^{\max} \in \hat{H}^R$ with almost certainty. When $p_t^{\min} = \max \{p \in H^A\}$ and $p_t^{\max} \in \hat{H}^R$, it follows that $p_t = p_t^{\min}$, which increases the likelihood that $p_{t+1}^{\min} = \max \{p \in H^A\}$. Therefore, as $t \rightarrow \infty$, $\Pr [p_t^{\min} = \max \{p \in H^A\}] \rightarrow 1$. From here, it follows that $p_{cip}^* = p_t^{\min} = \max \{p \in H^A\}$. ■

Proof of proposition 13. It is possible to observe any subset of all past transactions. Therefore, it is feasible that the seller observes the highest past accepted price and the lowest past rejected price during the first τ periods of the game. As τ increases, this becomes less likely, but is still possible. In this situation, the process converges just as it would under perfect information during the first τ periods. For τ large enough, the process will establish price p_{pc}^* prior to period τ . Since no long-term convention can be established at a price less than the highest past accepted price, it follows that p_{cip}^* will be at least as large as p_{pc}^* in this case. Therefore, $p_{cip}^* \geq p_{pc}^*$ is possible. ■