Seller preferences for risk seeking and limited information in an evolutionary price demand game

Christopher Cotton*

April 18, 2005

Abstract

The paper introduces evolutionary dynamics into a two-agent price demand game, in which sellers observe past period transactions before announcing a price, and buyers either accept or reject the announced price. Under the assumption of homogeneous clients, and large enough number of past-period observations, the process almost surely converges to a stable or continuously-reoccurring convention. Increased risk seeking in the class of sellers results in a long-term convention price at least as close to the buyer valuation as when the class of sellers is more risk averse. However, increasing risk seeking among sellers is often not feasible. I show that introducing imperfect information into the game by limiting the number of past period observations can also result in long-term convention prices closer to the buyer valuation, thereby increasing ex-ante expected utility among sellers. Thus sellers may have preferences for limited rather than perfect information of past transactions. However, buyers are never made better off by limiting seller memory.

---

*Preliminary and incomplete. I appreciate the comments of Ani Guerdjikova, Finn Christensen, and members of Cornell's grad student workshop in economics. All mistakes remain my own. Contact information: Department of Economics, Cornell University, Ithaca, NY; csc35@cornell.edu.

1 Introduction

Young's "An Evolutionary Model of Bargaining" (1993) applied evolutionary dynamics to a simple Nash-demand, or split-the-dollar game. He analyzed how the evolutionary process converges over time, and showed how the model favors certain divisions rather than others over the long run. His analysis relies on both players to simultaneously announce non-negotiable demands, and receive a positive payoff if and only if the sum of their demands does not exceed the total available surplus.[1] Therefore, his results are applicable to a variety of social interactions, but say very little about typical market interaction in which only one of the agents announces a price demand, and the other agent observes the announced price, and decides whether to accept or reject the price. This paper adapts Young's model such that members of only one agent class can announce a price demand, to develop a model that is more applicable to seller and buyer, or firm and client interaction. The analysis of this revised model suggests that firms/sellers may prefer less cooperation amongst the class of sellers where the result is limiting information or inducing risk-taking behavior.

This paper considers an infinitely repeated stage game in which members of two classes of agents are randomly selected and matched to play a price demand game. However, where Young assumed that both types of agents observed past transactions and announced prices, this model assumes that only members of one class of agents observe past-period transactions

[1] In Young's primary example, landlords and tenants present demands over shares of a crop. In each period, a landlord is randomly selected from a class of landlords, and a tenant is randomly selected from a class of tenants. The two agents observe a selection of landlord and tenant demands from previous periods, and then, given these observations, the agents develop beliefs regarding the likely demand announcements of their opponents given by a CDF. The density of the CDF is positive for a demand announcement if and only if the firm observed a past period in which a member of the other class announced that demand function. After determining their beliefs, they simultaneously announce their own demands x and y. They receive their respective demands if $x + y \leq 1$. If $x + y > 1$, both players receive nothing. After playing the game, the landlord and tenant exit the game, and are replaced in their respective classes by other agents of the exact same characteristics. This process is repeated over an infinite number of periods. Given this framework, Young develops an evolutionary bargaining process that describes how the process evolves between states, and shows how the process converges to conventions, which are states in which agents repeatedly announce demands x and y, where $x + y = 1$. Through his analysis, Young shows how, given sufficiently few observations, the process always converges to a convention, and that the achieved convention depends on the types of agents included in each class of agents, but not their relative proportions within their classes. He also shows how a convention involving 50-50 division results as the unique solution to the game under sufficient uncertainty, and mixing between classes.

and announce prices. Throughout most of my discussion, I assume that the price setting agents are sellers of a good or service. The selected member of the other class (consider the class of buyers) observe the seller's price demand and decides whether to accept or reject the price. If the buyer accepts, the seller receives payoff equal to the price, and the buyer receives a payoff equal to the difference between its valuation and the price. If the buyer rejects, both agents receive nothing. This stage game is repeated over an infinite horizon.

Given this framework, I develop an evolutionary price demand process that describes how the process evolves between states, and shows how the process converges to various types of conventions, or states in which sellers repeatedly announce the same price demand regardless of their type, and buyers always accept. After defining the evolutionary process, I consider how different assumptions regarding the level of seller risk aversion or risk seeking, and the number of past periods that sellers observe before announcing a price can impact the proximity of the convention price to the buyers' valuation. Although the model allows for the analysis of many different cases, this effort focuses on the situations where sellers observe transactions in all past periods of the game, and, alternatively, when their power of observation, or "memory," is limited to some finite number of periods. Additionally, the analysis focuses on the limiting case where the class of buyers is homogeneous, or sellers can perfectly observe buyer type, although the class of sellers can be either homogeneous or heterogeneous.[2]

Under fairly general conditions, the evolutionary process almost surely converges to a long-term convention. When sellers have perfect memory of past periods, the process achieves a stable convention, which, once established, remains established indefinitely. Alternatively, when seller observations are limited to a finite number of periods, the process achieves a continuously-reoccurring convention, where the process may occasionally deviate from, but always converge back to, the convention.

The price associated with these long-term conventions depend on the level of risk aversion,

[2] Although the same results hold for the case when buyer type is perfectly observed by sellers, the notation, discussion and proofs treat the class of buyers as homogeneous.

as well as the ex ante expectations regarding client valuations, within the class of sellers. Generally, the more risk seeking sellers are, the closer the long-term convention price is to the buyers' valuation. Additionally, limiting seller memory to a finite number of periods results in a long-term convention price at least as close to the buyers' valuation as is achieved when sellers have perfect information regarding past transactions.

From these results, it follows that under certain, entirely feasible conditions, aggregate seller earnings and utility can be improved by systems or policies that either limit seller observations, or promote risk seeking. When the sellers are thought of as sales people within a single firm, these results suggest possible motivation for the firm to promote competition between employees, such as providing the best performers with high bonuses while firing the worst performers. However, when increasing the level of employee risk aversion is not feasible, similar results may be achieved by limiting employee access to information. Additionally, when the the individual sellers are thought of as firms that are risk-neutral in profits, the ability to limit memory of past transactions becomes increasingly important.

The following section describes the model and defines the evolutionary process. The third section first defines a variety of convention types, then shows how the evolutionary process converges to these conventions under different assumptions regarding the level of information and the heterogeneity or homogeneity of the seller class. Given these results, the paper then describes the requirements for limited rather than perfect knowledge of past transactions to improve the sellers' ex-ante expected utilities. The final section presents a brief discussion regarding possible expansions of this paper.

2 Model[3]

There are two classes of individuals, A and B that are made up of a continuum of individuals. A represents the class of sellers of a good or service, and B represents the class of buyers of

[3]This model is directly built upon Young's evolutionary bargaining model. Therefore, in this section, the structure of certain paragraphs and sentences may be very similar or identical. These similarities help maintain an obvious relationship between this paper and the paper on which it is inspired.

that good or service. We can think of class A as a firm, and each agent $\alpha \in A$ as a sales person within that firm; or alternatively class A may represent a trade organization, and each $\alpha \in A$ may represent a firm with membership in that trade organization. The game has an infinite time horizon, and in each period $t = 1, 2, 3, ...$ one seller α_t is drawn at random from class A, and one buyer β_t is drawn at random from class B. α_t then observes some past periods of the game, then announces a price demand for the good or service, which we denote by p, and β_t chooses whether to accept or reject the price. If the buyer, β_t, accepts the price, the buyer pays amount p to the seller in exchange for the good or service. If the buyer rejects the price, the transaction does not take place, and both the buyer and seller receive nothing. This acceptance process is in contrast to Young's model, in which both players α and β announce demands, and received their demands if and only if they sum to no more than the total available amount. In this model, only the seller announces a price, and the buyer must decide whether or not to accept the price.

buyers may differ from each other in terms of their valuation of the good, as well as their utility functions. Let P_β denote β's valuation of the good, which also describes β's type; and $v_\beta(P_\beta - p)$ denote the utility that β receives from the transaction, such that $v_\beta(0) = 0$, and $v'_\beta(P_\beta - p) > 0$. Where required for clarity, I use the notation α_t to represent the seller that is selected to play in period t, and β_t to represent the buyer that is selected to play in the same period. Generally, sellers may have different utility functions $u_\alpha(p)$; although, for all sellers, $u'(p) > 0$. Sellers may also have different ex ante expectations of the buyers' valuations.

I do not require sellers to be risk neutral in earnings; allowing for both $u''(p) > 0$ and $u''(p) < 0$ to be possible. This is motivated in part because the sellers may be thought of as individuals rather than firms. It is also motivated by my desire to consider ex ante preferences to commit to policies or systems that promote risk seeking amongst the class of sellers.

For technical reasons, similar to Young's model, we assume that $p, P_\beta \in [P_{\min}, P_{\max}]$ and that there exists a finite number of feasible prices and valuations. Without loss of generality, we can normalize the range of prices such that $P_{\min} = 0$ and $P_{\max} = 1$; however, this is not required, and the majority of the paper is not discussed in these normalized terms. Let r be a positive integer, and let D be the set of all r-place decimal fractions that are positive and less than or equal to one. D is the set of feasible prices, and $\delta = 10^{-r}$ is the precision of the prices. This implies that $p, P_\beta \in \{P_{\min}, P_{\min} + \delta, ..., P_{\max} - \delta, P_{\max}\}$; however, through some abuse of notation, I refer to this discrete set of potential prices by the continuum $[P_{\min}, P_{\max}]$. These technical assumptions help assure that the process can achieve a convention in which the same prices are necessarily announced in sequential periods rather than only converge to prices arbitrarily close to the convention value.[4]

Let p_t denote the price demand 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 *demand history* for period t, and the sequence $(p_1, a_1), (p_2, a_2), ..., (p_t, a_t)$ denotes the *complete demand history* of the game up to and including period t.

Suppose agents α and β are chosen in period $t+1$. As with Young's model, we assume that neither party has prior knowledge or beliefs regarding the utility functions of the other agent, or about the distribution of the utility functions in the general population. However, unlike in Young's model, sellers have prior expectations regarding the distributions of the valuations held by members of class B. This additional assumption makes the process of updating beliefs more straightforward, while also making the model more realistic. Let the CDF $\bar{F}_\alpha(\cdot)$ represent α's ex ante beliefs regarding the possible distribution of the buyer's valuation P_β, and $\bar{f}_\alpha(\cdot)$ represent the distribution's density, such that $\bar{f}_\alpha(P) > 0$ for all $P \in [P_{\min}, P_{\max}]$. Therefore, when selected to play the game, sellers believe that any of the potential valuations on D are possible.

[4]Assuming a continuum of potential prices will not change the analysis results if we allow a convention at price p to be achieved when price announcements are necessarily within a neighborhood sufficiently close to p.

Before announcing a price, sellers observe the demand history from some past periods. Formally, we assume that α draws and observes a random sample of k_α of the past m records $((p_{t-m+1}, a_{t-m+1}), ..., (p_t, a_t))$. Therefore, the ratio k_α/m is a measure of α's *information*, and represents the extent to which α has knowledge of past periods. $K_{t(\alpha)}$ denotes the set of k_α histories observed by seller α from the last m records, and $t(\alpha)$ represents the period in which seller α is selected to play the game. Given these observations of past periods, the seller then updates its beliefs, forming a new CDF $F_\alpha\left(P \mid K_{t(\alpha)}\right)$ defined by the density function $f_\alpha\left(P \mid K_{t(\alpha)}\right)$. The seller then selects a price demand given $F_\alpha\left(P \mid K_{t(\alpha)}\right)$.

The structure of the process through which the seller updates its beliefs depends on whether the class of buyers is homogeneous or heterogeneous. Although the heterogenous case with multiple types of buyers is likely the more realistic case, we can still learn a lot from the suggestive case involving a homogeneous buyer class. To better identify the impact of differences in the class of sellers, my current effort focuses on a homogeneous class of buyers, and only provides a brief discussion of the alternative case, reserving expansion of the analysis for later work.

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

$$\Pr\{p \leq P_\beta\} = 1 - F_\alpha\left(p - \delta \mid K_{t(\alpha)}\right)$$

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

Therefore, seller α solves:

$$\max_{p \in D} u_\alpha(p)\left[1 - F_\alpha\left(p - \delta \mid K_{t(\alpha)}\right)\right]$$

The agents' response rules determine a stationary Markov chain. Let H_t represent the demand history of the game at time t, such that all possible sets of observations K_t are subsets of H_t. Let $\theta_\alpha(p \mid H_t)$ be the conditional probability that seller α announces price demand 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 is a best reply by α to some possible observation by α given history H_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 demand history as H_t up through time $t-1$, but also has an additional demand history for period t given by (p_t, a_t). Let $\pi(\alpha, \beta)$ be the probability that α and β and drawn to play against each other in any period. Every pair of agents has a positive probability of being drawn, though it is not necessarily the same probability for all pairs. If the process has history H_t at time t, then it has history H_{t+1} at time $t+1$ with probability

$$\Phi_{H_t H_{t+1}} = \sum_{\alpha \in A} \sum_{\beta \in B} \pi(\alpha, \beta)\, \theta_\alpha(p \mid H_t)$$

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 *evolutionary price demand process* with precision δ, memory m, information parameter k_α/m, and best reply distribution p_α.

This process can be simplified when the class of buyers is assumed to be homogeneous, as it is for the majority of the paper. Under this assumption,

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

Additionally, if both sellers and buyers come from homogeneous groups, $\Phi_{H_t H_{t+1}} = \theta_\alpha(p \mid H_t)$.

3 Analysis

3.1 Conventions

Similar to Young's paper, this analysis relies on the concept of conventions. However, although the definitions of conventions are similar between this paper and in Young's model, the differences in the models' framework mean that the concept must be redefined here.[5] This section defines the various concepts related to conventions that are used throughout the analysis. By introducing them together in the same section, I hope to better convey the relationship between them then if they were introduced separately throughout the paper.

Definition 1 *The process achieves a **common action** when there exists p^* such that $p^* = \arg\max_p u_\alpha(p)[1 - F_\alpha(p - \delta \mid K_t)]$ for all possible α and all possible observed histories K_t.*

Definition 2 *The process achieves a **repeated common action** at time $t+1$ when the process achieves a common action p^* at time t, and $p_t = p^*$ implies that $p_{t+1} = p^*$ with probability one.*

Definition 3 *A **convention** is a homogeneous set of sequential demand histories in which the process achieves and maintains a repeated common action for some number of periods.*

Definition 4 *A **stable convention** is achieved at time t when $p^* = \arg\max_p u_\alpha(p)[1 - F_\alpha(p - \delta \mid K_t)]$ for all possible α and K_t, and $p_t = p^*$ implies $p_s = p^*$ for all $s > t$.*

In other words, a stable convention is a convention which, once it is established, remains established through all future periods independent of which seller α is selected or which histories are observed in subsequent periods of the game.

When considering situations with limited memories, where m is finite, it is also necessary to consider temporary conventions. A *temporary convention* is different from a stable

[5] Young defines the 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 $0 < x < 1$."

convention, in that once established, a temporary convention only remains established for a finite number of periods before the process faces the potential for deviation. After a temporary convention is established for a certain number of periods, then the convention's price demand p^* ceases to be the optimal choice for all possible α and observed histories $K_{t(\alpha)}$. When this happens, the process is no longer in the convention, even if by chance the drawn α and $K_{t(\alpha)}$ result in p^* still being played. Temporary conventions can be classified as continuously-reoccurring and short-run conventions.

Definition 5 *A* ***convention of length n*** *is a convention that is maintained for n sequential periods until in the $n+1$th period there exists α such that p^* is no longer an optimal strategy given all possible K_t.*

Definition 6 *p^* is a* ***continuously-reoccurring convention*** *if there does not exist a stable convention, and there does exists a t (sufficiently large) such that after period t, p^* is the only convention that is ever achieved, and is achieved on a reoccurring basis.*

Both stable and continuously-reoccurring conventions can be referred to as *long-term conventions*. This is not to say that a continuously-reoccurring convention is the same as a stable convention. After a continuously-reoccurring convention is established, the process will occasionally leave the convention. However, the convention will always eventually be reestablished after any deviation.

Definition 7 *p^* is a* ***short-run convention*** *if it is established in some period t, and there exists an $s > t$ such that for all periods at least as large as s, the process never again achieves a convention p^*.*

3.2 Analysis under Homogeneous buyers

The analysis of this paper focuses on the situation where the group of buyers is homogeneous. The majority of this paper concentrates on various cases under this assumption. The seller recognizes that buyers are homogeneous; and that its buyer is exactly the same as the buyers that past sellers played against in past periods. Therefore, the seller knows that the buyer's valuation is less than the minimum rejected price, and greater than or equal to the maximum accepted price in past transactions. To formalize this process, let $P_A \subset K_{t(\alpha)}$ be the subset of all observed past-period prices that were accepted; and let $P_R \subset K_{t(\alpha)}$ be the subset of observed past-period prices that were rejected.[6] Then we can define $\check{p}_t = \max\{p \mid p \in P_A\}$ and $\hat{p}_t = \min\{p \mid p \in P_R\}$. If P_A is an empty set, then $\check{p}_t = P_{\min}$. If P_R is an empty set, then $\hat{p}_t = \delta + P_{\max}$.

When a seller updates its priors, given its observations of past period transactions:

$$f\left(P \mid \check{p}_t, \hat{p}_t\right) = \frac{\bar{f}\left(P\right)}{\bar{F}\left(\hat{p}-\delta\right) - \bar{F}\left(\check{p}-\delta\right)}$$

for all $P \in [\check{p}, \hat{p})$, and 0 otherwise.[7] It is not necessary that the seller reassign probabilities given this form; however, it is necessary that the seller assign a probability of zero that $P_\beta \notin [\check{p}, \hat{p})$. Also note that notation continues to be abused, where $[\check{p}, \hat{p})$ represents the discrete set $\{\check{p}, \check{p}+\delta, \check{p}+2\delta, ..., \hat{p}-\delta\}$.

Remember that the seller chooses p, and gets p iff $p \leq P_\beta$. $\check{p}$ is the only value in range $[\check{p}, \hat{p})$ such that the seller knows for sure that $p \leq P_\beta$. Furthermore, the seller is always better off choosing $\check{p}$ than any price demand less than $\check{p}$. The probability that other values are accepted is therefore determined by the CDF $F\left(P \mid \check{p}_t, \hat{p}_t\right)$, where the probability that a

[6] Don't think of the A in P_A as relating to the class of firms, also defined by A. I just noticed this potential for confusion, and will fix the issue at a later time.

[7] 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 $F\left(P\right)$ which is defined by the densities $f\left(P\right)$, such that

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

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 $P \leq p_n$. $F\left(P\right)$ may reasonably take another form, so long as it combines the seller's prior beliefs and observations into a new CDF.

price demand is accepted is given by:

$$\Pr\{p \leq P_\beta\} = 1 - F(p - \delta \mid \check{p}_t, \hat{p}_t)$$

The inclusion of δ results from the discrete case, since $\Pr\{p < P_\beta\} = 1 - F(p)$, and δ is the minimum possible increase in price demand/valuation.

Therefore, the seller α solves:

$$\max_{p \in D} u_\alpha(p) \left[1 - F_\alpha\left(p - \delta \mid \check{p}_{t(\alpha)}, \hat{p}_{t(\alpha)}\right)\right]$$

where $t(\alpha)$ denotes the period in which α is drawn to play.

Given this, the seller obviously chooses a value p such that $p \in [\check{p}, \hat{p})$. Where required for clarity, $\check{p}_t$ and $\hat{p}_t$ represent the values that $\check{p}$ and $\hat{p}$ take on in period t as defined above. The range of possible price demands and valuations in period t is given by $[\check{p}, \hat{p})$, and is called the *potential price range*. The potential price range is said to *converge* in period t if $\hat{p}_{t+1} - \check{p}_{t+1} < \hat{p}_t - \check{p}_t$, and $\check{p}_{t+1} \in [\check{p}_t, \hat{p}_t)$ and $\hat{p}_{t+1} \in (\check{p}_t, \hat{p}_t]$. Similarly, a process is said to *converge to a convention* if the potential price range converges to a range of values for which a convention is achieved given the other parameters of the model.

The process achieves a convention in period w, when for any possible α_w, the utility that the sellers gets for sure by choosing price demand $\check{p}_w$ must be at least as large as the expected utility from choosing a different potential price.

$$u_{\alpha_w}(\check{p}_w) \geq u_{\alpha_w}(\tilde{p}_w)\left[1 - F(\tilde{p}_w - \delta)\right] \text{ for any } \tilde{p}_w \in \{\check{p}_w, ..., \hat{p}_w - \delta\}$$

This can be rearranged to give the following convention condition:

$$\begin{aligned} u_{\alpha_w}(\tilde{p}_w) F(\tilde{p}_w - \delta) &\geq u_{\alpha_w}(\tilde{p}_w) - u_{\alpha_w}(\check{p}_w) \quad \text{or} \\ u_{\alpha_w}(\check{p}_w) F(\tilde{p}_w - \delta) &\geq \left[u_{\alpha_w}(\tilde{p}_w) - u_{\alpha_w}(\check{p}_w)\right]\left[1 - F(\tilde{p}_w - \delta)\right] \end{aligned}$$

where $u_{\alpha_w}(\tilde{p}_w) - u_{\alpha_w}(\check{p}_w)$ denotes the marginal utility from receiving price $\tilde{p}_w$ instead of price $\check{p}_w$. Therefore, $[u_{\alpha_w}(\tilde{p}_w) - u_{\alpha_w}(\check{p}_w)][1 - F(\tilde{p}_w - \delta)]$ is the expected benefit, and $u_{\alpha_w}(\check{p}_w) F(\tilde{p}_w - \delta)$ is the expected loss, from playing $\tilde{p}_w$ instead of price $\check{p}_w$. When the potential price range has converged to a situation where $\check{p}_w = \hat{p}_w - \delta$, then the potential price range is composed of a single value, $\tilde{p}_w \in \{\check{p}_w\}$, and the process achieves a convention. However, if $\check{p}_w < \hat{p}_w - \delta$, then whether the convention condition is satisfied depends on the functional form of $u_\alpha(\cdot)$ and $F(\cdot)$. For example, as a seller becomes increasingly risk averse, its choice of p will tend towards the lower end of the potential price range, and the seller becomes relatively more likely to favor $\check{p}$, the only value for which the seller expects its price demand to be accepted for sure.

Proposition 8 *For non-risk-seeking agents $\alpha \in A$, and $P_\beta > P_{\min}$, there exists some distribution of prior beliefs regarding P_β such that the process will converge to a convention p^* where $p^* < P_\beta$.*

I reserve the proof of this proposition for a future version of this paper. Although the situation will not always be such that the process converges to $p^* < P_\beta$, it is the case under some reasonable assumptions. The implication of this will become more obvious when I review the differences in models with full knowledge of history and limited memory in the following sections.

3.2.1 Full knowledge of history

To continue the analysis, we begin by considering the extreme case in which class A agents have knowledge of all past periods of the game. Technically, this implies that $k = m$ and m is infinitely large. However, we can simply think of the situation as sellers being perfectly informed regarding the history of the game. Young begins his analysis of the process in period m. Since full knowledge of past periods implies that period m is never achieved, this analysis begins with the first period.

In the first period, the seller α_1 chooses a price demand p_1 to solve:

$$\max_{p \in D} u_{\alpha_1}(p)\left[1 - F_{\alpha_1}(p - \delta \mid \check{p}_1, \hat{p}_1)\right]$$

Since there is no existing game history when the seller chooses its price demand, the seller cannot update its ex ante beliefs regarding buyer valuation; therefore $\check{p}_1 = P_{\min}$, $\hat{p}_1 = P_{\max} + \delta$, and $F_{\alpha_1}(\cdot \mid \check{p}_1, \hat{p}_1) = \bar{F}_{\alpha_1}(\cdot)$. Following the choice of p_1 by the seller, the first period buyer β_1 either accepts or rejects the price demand. If the price demand is accepted, α_1 receives p_1, and β_1 receives $P_\beta - p_1$.

In the second period, α_2 chooses a price demand p_2 to solve:

$$\max_{p \in D} u_{\alpha_2}(p)\left[1 - F_{\alpha_2}(p - \delta \mid \check{p}_2, \hat{p}_2)\right]$$

where $F_{\alpha_2}(\cdot)$ is no longer equivalent to the original CDF since the seller now observes p_1 and a_1, where $a_1 = 1$ if the buyer accepted the price demand, and $a_1 = 0$ if the buyer rejected the price demand.[8] Given these observations, $F_{\alpha_2}(\cdot)$ is updated given the process outlined above. If $a_1 = 1$, then $\check{p}_2 = p_1$, and $\hat{p}_2 = P_{\max}$; which implies that $f_2(P \mid \check{p}_2, \hat{p}_2) > 0$ for all $P \in [p_1, P_{\max}]$ and $f_2(P \mid \check{p}_2, \hat{p}_2) = 0$ for all $P \in [P_{\min}, p_1)$. Alternatively, if $a_1 = 0$, then then $\check{p}_2 = P_{\min}$, and $\hat{p}_2 = p_1$; which implies that $f_2(P \mid \check{p}_2, \hat{p}_2) > 0$ for all $P \in [P_{\min}, p_1)$ and $f_2(P \mid \check{p}_2, \hat{p}_2) = 0$ for all $P \in [p_1, P_{\max}]$. Given these updated beliefs, if p_1 was accepted, α_2 will always choose $p_2 \in [p_1, P_{\max}]$; and if p_1 was rejected, α_2 will always choose $p_2 \in [P_{\min}, p_1)$.

In the third period, the randomly selected seller α_3 observes the play from the previous two periods, which is given by history $((p_1, a_1), (p_2, a_2))$. We can therefore consider α_3's reaction given the four possible histories at this time: $((p_1, 1)(p_2, 1))$; $((p_1, 1)(p_2, 0))$; $((p_1, 0)(p_2, 1))$; and $((p_1, 0)(p_2, 0))$. Remember a_t is the indicator variable describing whether p_t was accepted by the buyer in period t.

When the history is $((p_1, 1)(p_2, 1))$, both of the previous-period offers were accepted,

[8] $F_2(\cdot)$ does not equal the original CDF unless $p_1 = P_{\max}$.

and therefore α_3 knows that the buyer's valuation P_β is greater than or equal to both of the previous price demands, p_1 and p_2. Since it is also clear that $p_2 \in [p_1, P_{\max}]$, we know $\check{p}_3 = p_2$ and $\hat{p}_3$ continues to equal $P_{\max}$; and therefore α_3 chooses $p_3 \in [p_2, P_{\max}]$. If $p_3 \neq p_2$, or rather $p_3 \in (p_2, P_{\max}]$, then the potential price range converges in period 3. However, if the optimal choice of $p_3 = p_2$, then $\check{p}_3 = \check{p}_2$, and α_4 will have no additional knowledge compared to α_3; therefore the potential price range does not converge.

Alternatively, if the history is $((p_1, 1)\,(p_2, 0))$, the first period offer was accepted, while the second period offer was rejected. Therefore, α_3 can conclude that the consumers' valuation is at least as large as p_1 and less than p_2, which means that $\check{p}_3 = p_1$ and $\hat{p}_3 = p_2$. It will therefore choose $p_3 \in [p_1, p_2)$. The third case with history $((p_1, 0)\,(p_2, 1))$ is opposite the second case. In this situation, α_3 concludes that the consumers' valuation is at least as large as p_2 and less than p_1, which means that $\check{p}_3 = p_2$ and $\hat{p}_3 = p_1$. She will therefore choose $p_3 \in [p_2, p_1)$. In both of these cases, if $p_3 = \check{p}_3$, then no additional information is provided to α_4, and the potential price range does not converge. However, if p_3 is a new price demand, such that $p_3 \in [p_2, p_1)$, then the potential price range does converge.

In the final case, the history is $((p_1, 0)\,(p_2, 0))$, and α_3 observes that both of the previous offers were rejected. She therefore concludes that the consumers' valuation P_β must be strictly less than both of these previous-period price demands. With similar reasoning to the first case, since $p_2 \in [P_{\min}, p_1)$, it follows that $\hat{p}_3 = p_2$. The knowledge regarding the buyers' valuation can then be described by $P_\beta \in [P_{\min}, p_2)$. Therefore, $p_3 \in [P_{\min}, p_2)$. Unlike with the other three cases, p_3 cannot equal either of the previous price demands, and therefore, the potential price range will converge so long the seller chooses a price demand greater than $P_{\min}$.

In later periods of the game, the process behaves similarly. It can be seen that following each period, as an additional period of history is observed by future sellers, the range of possible valuations will either decrease, or will remain unchanged. The range decreases between period t and period $t+1$ when the value p_t does not equal any of the price demands from previous periods; in which case it is a new price demand between values $\check{p}_t$ and $\hat{p}_t$. However, if p_t equals a price demand from a previous period the same information will be available for the seller in period $t+1$ as was available in period t; this only happens when $p_t = \check{p}_t$. With full knowledge of history, the potential price range either contracts or remains constant as the game progresses. Therefore, a seller will always have at least as much information, or, more specifically, at least as small of range for possible prices, as any of its predecessors.

In any period t of the game, the seller assigns values to $\check{p}_t$ and $\hat{p}_t$ as described above, and effectively ignores all other past demands. By definition of $\check{p}_t$ and $\hat{p}_t$, with full knowledge of past histories there cannot exist a previous period demand between these values, therefore $p_s \notin (\check{p}_t, \hat{p}_t)$ for all $s \leq t$. The seller will never choose a value $p_t < \check{p}_t$, since the seller recognizes that a higher price demand $p_t = \check{p}_t$ will be accepted with probability 1; and will never choose a value $p_t \geq \hat{p}_t$ since the seller recognizes that such a price demand will be rejected with probability 1. Therefore, to maximize $u_{\alpha_2}(p)\left[1 - F_{\alpha_2}(p - \delta \mid \check{p}_2, \hat{p}_2)\right]$, the seller will always choose a price demand p_t such that $p_t \in [\check{p}_t, \hat{p}_t)$.

The price demand p_t is either one of two types. We call p_t a *new price demand* if it was not announced in an observed period, which is the characteristic of any $p_t \in (\check{p}_t, \hat{p}_t)$. The new price demand may be accepted or rejected. If the price demand is accepted, then the seller in the next period updates its beliefs such that $\check{p}_{t+1} = p_t$, and $\hat{p}_{t+1} = \hat{p}_t$. If the price demand is rejected, then the seller in the next period updates its beliefs such that $\check{p}_{t+1} = \check{p}_t$ and $\hat{p}_{t+1} = p_t$. Alternatively, p_t is called a *repeated price demand* if it was announced as a price demand in an observed period. With homogeneous buyers, this can only be true when $p_t = \check{p}_t$. If α_t announces a repeated price demand, then the seller in the next period, α_{t+1}, will have the same information as α_t, such that $\check{p}_{t+1} = \check{p}_t$ and $\hat{p}_{t+1} = \hat{p}_t$; and therefore the potential price range does not converge. Notice that with full information, the potential price range converges in period t if and only if p_t is a new price demand.

Proposition 9 *When sellers have perfect knowledge of the price demand history, and the* *class of buyers is homogeneous, the process will always achieve a stable convention.*

Proof. Claim: The potential price range can converge at most $(P_{\max} - P_{\min}) \frac{1}{\delta}$ times over the entire course of the game. Subproof: $(P_{\max} - P_{\min} + \delta) \frac{1}{\delta}$ is the total number of feasible price demands and valuations. Even with minimal convergence each time, the process will achieve a situation where $\check{p}_t = \hat{p}_t - \delta$ after the potential price range converging at most $(P_{\max} - P_{\min}) \frac{1}{\delta}$ times.

Claim: When $\check{p}_t = \hat{p}_t - \delta$, the process is in a convention. Subproof: As we established above through the definition of $\check{p}_t$ and $\hat{p}_t$, the seller always chooses $p_t \in [\check{p}_t, \hat{p}_t)$, which is equivalent to $p_t \in [\check{p}_t, \hat{p}_t - \delta]$. Therefore, when $\check{p}_t = \hat{p}_t - \delta$, $p_t \in \{\check{p}_t\}$, thus $\check{p}_t = \arg\max_p u_\alpha(p) \left[1 - F_\alpha\left(p - \delta \mid \check{p}_{t(\alpha)}, \hat{p}_{t(\alpha)}\right)\right]$ for any $\alpha \in A$. K_α is independent of α, and there is no memory loss. Thus, $p^* = \check{p}_t$.

Now, consider the behavior of seller in time t. α_t chooses $p_t \in [\check{p}_t, \hat{p}_t)$. If $p_t \in (\check{p}_t, \hat{p}_t)$, the potential price range converges. Alternatively, if $p_t = \check{p}_t$, then the potential price range does not converge. In this subcase, if $\check{p}_t$ is an optimal choice for all sellers, then the process achieves a convention. However, if $\check{p}_t$ is not an optimal choice for all sellers, then the process will eventually draw an α that selects a price demand between $\check{p}_t$ and $\hat{p}_t$ causing the potential price range to converge. Since the potential price range can converge at most $(P_{\max} - P_{\min}) \frac{1}{\delta}$ times, if the process does not achieve a convention beforehand and stop converging, it will achieve one when the convergence results in $\check{p} = \hat{p} - \delta$.

Claim: If the process achieves a convention at time t, then $p^* = \check{p}_t$. Subproof: Straightforward.

Claim: With full information of the complete history, any convention is a stable convention. Subproof: By definition, in a convention, all types of sellers have the same optimal value p^* for any possible information draw. Since $p^* = \check{p}_t$, no additional information is created when the process is in a convention; and since there is full information of history, no information is loss. Therefore, p^* continues to be the optimal choice for all sellers as the game progresses, and is a stable convention.

This concludes the proof.[9] ■

It is also interesting to consider differences in the processes through which homogeneous and heterogeneous sellers achieve a convention. In the case of homogeneous sellers, any achieved convention in which $\check{p}_t = \arg\max_p u_\alpha(p)\left[1 - F_\alpha\left(p - \delta \mid \check{p}_t, \hat{p}_t\right)\right]$ is a stable convention. The same information is available to subsequent sellers, and therefore these other sellers from the homogeneous class of sellers act exactly the same as the seller in period t. This implies that when the class of sellers is homogeneous, the actions of the sellers are independent of which seller is selected, and the impact that random draws of nature may have are completely eliminated from the model. Therefore, the convention that is achieved, and the period in which it is achieved are independent of the sequence of sellers that are

[9] Achieving a convention is therefore independent of whether the class of sellers is homogeneous or heterogeneous. However, to provide intuition as to this proposition, I consider the game with a homogeneous group of sellers, then show that the same conclusions hold when the class of sellers is heterogeneous.

First, considering the case of homogeneous sellers. Since sellers have the same priors, it holds that if $\hat{p}_t = \hat{p}_r$ and $\check{p}_t = \check{p}_r$, then $F_\alpha(\cdot \mid \check{p}_t, \hat{p}_t) = F_\alpha(\cdot \mid \check{p}_r, \hat{p}_r)$. Therefore, if $\check{p}_t = \arg\max_p u_\alpha(p)\left[1 - F_\alpha\left(p - \delta \mid \check{p}_t, \hat{p}_t\right)\right]$, then α_t plays the repeated price demand, providing no new information to the seller in period $t+1$. It follows that $\hat{p}_t = \hat{p}_{t+1}$ and $\check{p}_t = \check{p}_{t+1}$, and thus $F_\alpha(\cdot \mid \check{p}_t, \hat{p}_t) = F_\alpha(\cdot \mid \check{p}_{t+1}, \hat{p}_{t+1})$. Following this through, we see that $\check{p}_t = \arg\max_p u_\alpha(p)\left[1 - F_\alpha\left(p - \delta \mid \check{p}_{t+1}, \hat{p}_{t+1}\right)\right]$ as well, and therefore, sellers in both period t and $t+1$ announce the same price demands. It is straightforward to extend this logic to consider any period after t. As in period $t+1$, the seller selected in period $t+2$ will have no additional information compared to α_t, and will therefore choose the same price demand as α_t and α_{t+1}. The reason behind this is the same reason why α_{t+1} chooses the same price demand as α_t, which I discussed above. This logic continues from period t through all future periods. Therefore, in any period $s > t$, $p_s = p_t = \check{p}_t$.

Although this argument shows that for the case of homogeneous sellers, a seller's choice in to play $\check{p}_t$ in period t is sufficient to assure that sellers in all periods following t will also choose to play $\check{p}_t$, the argument in itself does not guarantee that there exists a period t such that $\check{p}_t = \arg\max_p u_\alpha(p)\left[1 - F_\alpha\left(p - \delta \mid \check{p}_t, \hat{p}_t\right)\right]$. To see why this situation is assured, we must remember that there are a finite number of possible price demands (as determined by the precision of the range δ, described above), and that in each period, the process either converges, or remains constant. If the process remains constant between two periods t and $t+1$, the seller must have played $p_t = \check{p}_t$. Otherwise, if the seller played $p_t \in (\check{p}_t, \hat{p}_t)$, new information would be available, and the process would converge. If on the other hand the process does converge, since $p_t \in (\check{p}_t, \hat{p}_t)$, then the seller in the following period has fewer price demand choices than the seller in period t. If this is continuously repeated, with sellers always preferring values for p other than $\check{p}$, the process will eventually reach a situation in which $\hat{p} = \check{p} + \delta$. When this happens, $\check{p}$ is the only possible value for P_β, and thus the only rational price demand. Therefore, in an infinite horizon game, the process will always converge to a convention.

This conclusion also holds true when the group of sellers is not homogeneous. As suggested in the argument for homogenous sellers, the process cannot converge more times than there are possible price choices. This characteristic is independent of whether the class of sellers is homogeneous or heterogeneous. Therefore, the maximum number of times that a process can converge before achieving a convention is equal to the number of possible prices. More precisely, the process can converge in at most $(P_{\max} - P_{\min})\frac{1}{\delta}$ periods. As with the case of homogeneous sellers, the process likely achieves a convention before this number of periods. This is further outlined in the more formal proof above.

drawn to play the game. With homogeneous sellers and buyers, and perfect information of history, it can be determined with certainty the convention that the process achieves, and the period in which the convention is achieved. This idea is summarized in the following proposition:

Proposition 10 *When the classes of sellers and buyers are homogeneous, and sellers have perfect information regarding the history of the game, there exists some convention p^* and period w, such that the process achieves stable convention p^* in period w with probability one.*

The same result does not necessarily hold with a heterogeneous class of sellers. Although the process will eventually achieve a continuous convention, uncertainty regarding which seller type is drawn to play the game in each period prevents a similar proposition applying when the class of sellers is heterogeneous. Even if at period t, $\check{p}_t = \arg\max_p u_{\alpha_t}(p)\left[1 - F_{\alpha_t}(p - \delta \mid \check{p}_t, \hat{p}_t)\right]$, there may exist a following period $s > t$, such that a_s is a different type of seller than α_t, which causes either (or both) the priors regarding the buyers' valuation or the utility functions to differ. This would mean that the two sellers are actually solving different utility maximization problems since

$$u_{\alpha_t}(p)\left[1 - F_{\alpha_t}(p - \delta \mid \check{p}_t, \hat{p}_t)\right] \neq u_{\alpha_s}(p)\left[1 - F_{\alpha_s}(p - \delta \mid \check{p}_t, \hat{p}_t)\right]$$

This may mean that $\check{p}_t \neq \arg\max_p u_{\alpha_s}(p)\left[1 - F_{\alpha_t}(p - \delta \mid \check{p}_t, \hat{p}_t)\right]$. Therefore, α_s plays a new price demand, causing the potential price range to further converge. As discussed above, this can only happen at most $(P_{\max} - P_{\min})\frac{1}{\delta}$ times, and the process will still eventually achieve a convention. However, with heterogeneous sellers, the convention may take a greater number of periods to achieve, since convergence in every period until a convention is achieved is not assured.[10]

Additionally, further consideration of the situation suggests some interesting results. We can conclude that the more risk seeking sellers are, the more likely the process is to achieve

[10] I might want to eventually discuss the probability that the process achieves convention p^* in period w by determining a formula describing these values for any p^* and w.

a convention relatively close to the buyers' valuation. Similarly, the more risk averse the sellers, the more likely the process achieves a convention significantly less than the buyers' valuation. Also, and more obviously, the more accurate the seller's prior beliefs are, the more likely the process converges to a convention near the buyer's valuation. A heterogeneous group of sellers, does not guarantee the eventual convergence to a convention that is closer to the buyers' actual valuation than would be achieved with a homogeneous group of sellers; however, it is likely that greater diversity within the class of sellers creates the potential for faster convergence, and convergence to a convention closer to P_β. Whether or not this holds depends on the specific characteristics of the seller types.[11]

3.2.2 Limited memory

The previous section considered the model when sellers have perfect knowledge of the game's history. This section weakens this assumption, and considers a case where sellers only have knowledge of the most recent m periods. Similar to the last section, we continue to assume that technically $k = m$; however, m is now finite. Because I am 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. Any $m \geq \frac{P_{\max}-P_{\min}}{\delta}$ is always sufficiently large enough 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. This section begins by considering the case of homogeneous sellers, and draws conclusions regarding the convergence of the process when this assumption holds. I then consider how the results are altered if the class of sellers is heterogeneous.

Again, unlike in Young's approach, I do not begin the analysis in period m, but rather discuss the process as beginning in the first period. This assumption has less justification than in the previous section, since m is no longer infinite. However, any conclusions from

[11] These points will be more formally addressed when this paper is further developed.

an analysis starting in the first period will also hold if the analysis started by considering a random mth period. Up until period m, sellers observe all past periods of the game. After period m, sellers observe only the most recent m periods of the game. Therefore, between the first period and period m, the game proceeds exactly as it did under the assumption of full information. This will either result in a convention being achieved prior to period m, or movement of the process towards a convention that was not achieved by period m. To consider the implications of having an m-period memory, I begin by considering the impact of a memory limit on a convention once it is achieved. Initially, let m be large enough such that, when the class of sellers is homogeneous, the process achieves an initial convention in period $w < m$.

As the previous section showed, with full knowledge of history, the process always achieves a convention in some finite number of periods. When the class of sellers is homogeneous, uncertainty about the type of seller drawn to play in any period is eliminated, and when sellers have perfect knowledge of the games history, the process always converges to the same convention in the same period. Therefore, if we assume that seller memory is longer than it takes for the process to achieve a convention under full knowledge, the seller will continue to achieve the convention in the same period as it did under full knowledge, $w < m$. Therefore, a sufficiently long memory assures that the process converges to a convention with probability one. However, it is also important to recognize that $m \leq w$ may also achieve a convention in the game.

In the situation where $w < m$, the first convention that the process achieves is the same as the stable convention achieved under full knowledge of history. Therefore, it is just as likely as in the previous case that the convention price $p^* < P_\beta$. However, with limited memory, this convention is generally a temporary convention rather than a stable convention. This implies that after some period of time, the process leaves the convention, and potentially converges to a different convention. To see this, consider the following example.

Assume that the process achieves a convention p_w^* in period $w < m$, when the potential price range has converged to the discrete set of possible prices within $[\check{p}_w, \hat{p}_w)$. By definition, $p_w^* = \check{p}_w = \arg\max u_\alpha(p)[1 - F_\alpha(p - \delta \mid \check{p}_w, \hat{p}_w)]$ for all $\alpha \in A$. If $\check{p}_w = \arg\max u_\alpha(p)[1 - F_\alpha(p - \delta \mid \check{p}_w, P_{\max})]$, then the convention is stable, since, as I show below, limited memory results in $\hat{p}_t = P_{\max}$ at some future period $t > w$. However, generally $P_{\max} \neq \hat{p}_w$ and $p_w^* = \check{p}_w \neq \arg\max u_\alpha(p)[1 - F_\alpha(p - \delta \mid \check{p}_w, P_{\max})]$ for all $\alpha \in A$. . When this is the case, the convention is temporary.

Suppose the value $\hat{p}_w$, the minimum observed rejected price demand, corresponds to the announced price demand in period r such that $r \in \{w - m, ..., w - 1\}$. In all periods t in which the seller observes period r's rejected price demand, α_t will never announce a price demand $p_t \geq p_r$. Since $\hat{p}_w = p_r$, it also holds that $\hat{p}_t = p_r$ for all $t = r + 1, ..., w$. Additionally, given this framework, in periods $t = r + 1, ..., w$, p_t is always accepted, and is therefore less than or equal to $\check{p}_w$. Since a convention is achieved in period w, so long as the same information is available in other periods, with the sellers continuing to observe $\check{p}_w$ and $\hat{p}_w$, the convention is maintained. Since $\check{p}_w$ is played as the common action in the convention p_w^*, its value will remain in memory for at least m period after the convention ends. Therefore, forgetting the value $\check{p}_w$ will never be the reason that the process leaves a convention. However, the value $\hat{p}_t$ will be forgotten m periods after it was originally played. Given the above structure, since it was originally played in period r, it is no longer observed by the seller selected to play in period $r + m$. When this happens, the upper bound on the set of potential prices becomes $\hat{p}_{r+m} = P_{\max}$. To see this, remember that in all periods between r and $r + m$, sellers observed rejected p_r and therefore would not announce a price greater than it. When it is no longer observed, sellers also do not observe any other rejected prices, and therefore the maximum potential price returns to the ex ante default maximum price $P_{\max}$.

If $p_w^* = \check{p}_w \neq \arg\max u_\alpha(p)[1 - F_\alpha(p - \delta \mid \check{p}_w, P_{\max})]$, then α_{r+m} selects a price $p_{r+m} > p_w^*$. When this happens, the process will again converge towards a convention. If, as the process is again converging to a new convention, a seller announces price demand p such that $p_w^* < p \leq P_\beta$, the process converges to a new convention that is closer to the buyer's valuation than p_w^*. Even if the process does not converge to a new convention closer to P_β than the original convention, with sufficiently large m, sellers will never announce prices that are further from the buyers' valuation than p_w^*. Therefore, for all conventions p_t^*, such that $t > w$, and sufficiently long memory, it must be the case that $p_w^* \leq p_t^* \leq P_\beta$. If however $p_w^* = P_\beta$, then all future conventions must also achieve a price equal to P_β.

Proposition 11 *When the class of buyers is homogeneous, there exists a value s such that for any $m > s$ the process achieves a continuously-reoccurring convention with probability one.*

Proposition 12 *The price demand associated with the continuously reoccurring convention is at least as large as the price demand association with the stable convention that would have resulted under full knowledge of game history.*

Proof. Sketch.[12] First, consider the case of homogeneous sellers. For sufficiently large m, the same argument assures reconvergence to a convention as assured convergence to the initial convention. I go through this argument in the proof to an earlier proposition. Additionally, as I discuss above, for all conventions p_t^*, such that $t > w$, and sufficiently long memory, it must be the case that $p_w^* \leq p_t^* \leq P_\beta$. Let the convention following convention p_w^* be established in period t. If $p_t^* = p_w^*$, given the the class of sellers is homogeneous, and memory is sufficiently long, the process will always converge back to the same convention, and therefore convention $p_t^* = p_w^*$ is a continuously reoccurring convention. However, if $p_t^* > p_w^*$, then the process converges to a new convention. Then, after the process deviates from convention p_t^* it will again either reconverge to p_t^* or some other value greater than p_t^* and no larger than P_β. Since the set of potential prices is finite, the process will eventually achieve a convention $p^* = P_\beta$, or a convention $p^* < P_\beta$, but where the process always reconverges to p^*; in either case, a continuously-reoccurring convention.

[12] A more formal, carefully layed out proof will be provided in a future version of this paper.

When the class of sellers is heterogeneous, a similar argument assures that a continuously-reoccurring convention is eventually achieved. With sufficiently large m, the process will always eventually achieve a situation where an action is repeated in two consecutive periods, t and $t+1$. This implies that $p_t = p_{t+1} = \check{p}_{t+1} = \check{p}_{t+2}$. When this happens, one of two situations must hold true. Either the process is in a convention, or the process is not in a convention, but α_t and α_{t+1} found action $\check{p}_{t+1}$ optimal. Therefore, drawing a seller of type α_{t+1} in period $t+2$ assures that the action will be repeated again; although other types of $\alpha \in A$ may find the repeated action optimal as well. However, eventually, the process will generally deviate from this repeated action. This happens when either $\check{p}_{t+1} = \check{p}_t$ is no longer observed; or when an α is drawn such that $\check{p}_{t+1} \neq \arg\max u_\alpha(p)\left[1 - F_\alpha\left(p - \delta \mid \check{p}_{t+1}, \hat{p}_{t+1}\right)\right]$, noting that $\check{p}_{t+1} = \check{p}_{t(\alpha)}$ and $\hat{p}_{t+1} = \hat{p}_{t(\alpha)}$. When this happens, $p_{t(\alpha)-1} = \check{p}_{t+1}$, which implies that the lower bound of the potential price range is not forgotten for at least m additional periods; and therefore, with sufficiently large m, the process will achieve another repeated action at a price at least as large as $\check{p}_{t+1}$. Therefore, after a repeated price is achieved in period $t+1$, future repeated prices will be in the range $[p_{t+1}, P_\beta]$.[13] Similar to the previous case, the process then eventually achieves a state where it is in a continuously reoccurring convention. This follows since, if the process does not achieve one previously, it will continue to converge to repeated price demands at higher values of p, and will eventually achieve the state where $\check{p} = \hat{p} - \delta = P_\beta$. When this happens, the process always achieves a continuously-reoccurring convention. ■

3.3 Benefits of risk seeking and imperfect information

3.3.1 Benefit of seller risk seeking

As illustrated above, under the assumption of homogeneous buyers and long enough memory, the evolutionary price demand process almost surely converges to a long-term convention. With perfect information, a stable convention is achieved; and with limited memory, a

[13] Again noting the abuse of notation, where I use a continuum to denote the potential discrete values.

continuously-reoccurring convention is achieved. The long-term convention price depends on the ex ante expectations that firms have over the buyer's valuations $\bar{F}_\alpha(\cdot)$, and the form of the sellers' utility functions $u_\alpha(\cdot)$.

Proposition 13 *Holding ex ante expectations regarding the distribution of the buyers' valuation and the level of risk aversion of other seller types constant, decreasing the risk aversion of one seller type results in a long-term convention price at least as close to the buyers' valuation as when the seller type is more risk averse.*

A formal proof is reserved for a future version of this paper; however, intuition is provided here. Let $p^* < P_\beta$ denote the price associated with the long-term convention that is established from some arbitrary $\left\{\bar{F}_\alpha(\cdot), u_\alpha(\cdot)\right\}_{\alpha\in A}$ such that $\forall\alpha,\ u'_\alpha(\cdot) > 0$ and $\bar{f}_\alpha(P) > 0$ for all $P \in [P_{\min}, P_{\max}]$. Note that these are the standard assumptions regarding u and $\bar{F}$ discussed previously. Making at least one of the sellers more risk seeking can result in the seller choosing an alternative $p \neq p^*$ when presented with the original convention's potential price range $[p^*, \hat{p}^*)$. This is because firms that are more risk seeking may be more likely to risk getting nothing for the chance of getting a higher payoff rather thanjust get p^* for sure. This can result in the process converging to a long-term convention associated with a price $p' \in (p^*, P_\beta]$ rather than the original long-term convention price p^*.

Formally, holding ex ante expectations over buyer valuations, $\bar{F}_\alpha(\cdot)$, constant for all sellers, the process will converge to a convention price $p' \in (p^*, P_\beta]$ if at least one seller's risk aversion is decreased such that:

$$u_\alpha(p^* + \delta)\left[1 - F_\alpha(p^* \mid p^*, \hat{p}^*)\right] > u_\alpha(p^*)$$

where $[p^*, \hat{p}^*)$ represents the potential price range maintained in the original p^* convention. This condition always holds whenever α prefers to play some $p' \in (p^*, P_\beta]$ when observing $[p^*, \hat{p}^*)$.

Consider the example where class A represents a firm, and $\alpha \in A$ are the members of the firm's sales staff. In this case, firm A's profit may be risk neutral in aggregate long-term seller earnings. However, each $\alpha \in A$ need not be risk neutral in earnings. Firm utility may be denoted by it's revenue function $\sum_{t=1}^{T} p_t a_t$ or alternatively by its average per period revenue function:

$$U_A = \frac{1}{T} \sum_{t=1}^{T} p_t a_t$$

where T is the number of periods.

In the case of perfect information, as $T \to \infty$, $U_A \to p^*$ where p^* is the price associated with the long-term stable convention. If the firm can increase the level of risk seeking among at least some of its sales personnel to achieve a long-term convention price of $p' > p^*$, then it can improve its long-term revenue. A very similar result holds when information is limited to the last m periods. With limited memory, the proportion of periods in which the process is in a continuously reoccurring convention converges to a constant $\lambda \in (0, 1]$ as $T \to \infty$. Therefore, as $T \to \infty$, $U_A \to \lambda p^*$. Again, long-term revenue can be increased by increasing the risk-seeking behavior among its sales staff. These results suggest possible motivation for the firm to promote competition between employees, such as providing the best performers with high bonuses while firing the worst performers.

Additionally, aggregate seller utility may also be improved by increasing the risk taking behavior of all sellers. This also implies that, because the order each agent is selected to play the game is randomly chosen, the ex ante expected utility of individual sellers may also be improved. Although sellers may never prefer to artificially increase only their own risk seeking behavior, they may prefer a policy or system that encourages all sellers to take risks because they benefit from increases to information that can result from other players taking more risks.

3.3.2 Benefit of imperfect information

The price demand in the continuously-reoccurring convention achieved under finite memory is at least as close to the buyers' valuation P_β as the price in the stable convention achieved under full knowledge of game history. Because of this, limiting memory can have the same impact as increased risk seeking on the long-term earning of the class of sellers, and expected utilities of the sellers themselves. This becomes especially important when level of seller risk aversion cannot be influenced.

Continuing from the discussion of seller risk seeking, let $p^* < P_\beta$ denote the price associated with the long-term convention that is established from some arbitrary $\left\{\bar{F}_\alpha(\cdot), u_\alpha(\cdot)\right\}_{\alpha\in A}$ such that $\forall\alpha,\ u'_\alpha(\cdot) > 0$ and $\bar{f}_\alpha(P) > 0$ for all $P \in [P_{\min}, P_{\max}]$. Again, these are the standard assumptions regarding u and $\bar{F}$ discussed previously. Here, we now hold both $\bar{F}_\alpha(\cdot)$ and $u_\alpha(\cdot)$ constant for all α, thereby assuming no changes to either the ex ante beliefs over P_β or the levels of seller risk aversion. Instead, only the potential price range is allowed to change.

Long-term average seller utility is given by:

$$LTASU = \lim_{T\to\infty} \sum_{t=1}^{T} u_{\alpha_t}(p_t)\, a_t \frac{1}{T}$$

In the case where sellers have perfect knowledge of the game history:

$$LTASU_{FullKnowledge} = \sum_{\alpha\in A} \pi(\alpha)\, u_\alpha(p^*)$$

where p^* is the convention achieved in period w of the game. Remember that $\pi(\alpha)$ denotes the probability that a seller of type α is drawn from the class of sellers A in any given period. Also, once p^* is achieved under full knowledge, it remains established through the remainder of the game; any achieved convention under full knowledge is a stable convention.

Alternatively, suppose sellers only observe the most recent m periods of the game, where

m remains sufficiently large to achieve a convention. With limited memory, the proportion of periods in which the process is in a continuously-reoccurring convention converges to a constant $\lambda \in (0, 1]$ as $T \to \infty$. Let p' denote the long-term price associated with the continuously-reoccurring convention. As shown earlier, $p' \geq p^*$. Then,

$$\begin{aligned} LTASU_{LimitedMemory} &= (1-\lambda)0 + \lambda \sum_{\alpha \in A} \pi(\alpha) u_\alpha(p') \\ &= \lambda \sum_{\alpha \in A} \pi(\alpha) u_\alpha(p') \end{aligned}$$

Over the long run, when the process is not in a continuously-reoccurring convention, sellers are announcing prices that are rejected by the sellers, and therefore resulting in zero payoff. Otherwise, the process would continue to converge and achieve a different long-term convention than p'.

The long term average seller utility, and therefore the long-term aggregate seller utility, is only improved if:

$$\lambda \sum_{\alpha \in A} \pi(\alpha) u_\alpha(p') > \sum_{\alpha \in A} \pi(\alpha) u_\alpha(p^*)$$

which is clearly possible. When this condition holds, sellers can achieve an aggregate expected utility improvement by limiting memory.

Additionally, limited memory improves the ex ante expected utility of individual agent α when:

$$\lambda u_\alpha(p') > u_\alpha(p^*)$$

Therefore, it is clearly possible for all or some of the agents to prefer a system of limited rather than perfect information prior to the beginning of the game.

If again we alternatively consider the example where class A represents a profit maximizing firm, and $\alpha \in A$ are firm employees, the firm may decide to limit the access it's sales team has to information regarding past transactions. In this case, the required condition is the same as when we considered the implications of risk seeking behavior above. The firm's average per-period revenue function is given by

$$U_A = \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} p_t a_t$$

In the case of perfect information, as $T \to \infty$, $U_A \to p^*$; and in the case of limited information, as $T \to \infty$, $U_A \to \lambda p'$. Therefore, the firm prefers to limit employee access to information regarding past transactions when $\lambda p' > p^*$, which again is clearly possible.

Buyer welfare on the other hand, is never improved by limiting the memory of the seller. With sufficiently large m, the process converges to an initial convention p^*. Under perfect knowledge of the game history, the convention p^* is maintained as a stable convention. Alternatively, under limited memory, generally, the process will eventually deviate from the convention p^*, and then again converge to a convention. In the new convention, sellers are at least as well off as in the original convention, and buyers are no better off than in the original convention. Additionally, the possible loss in welfare due to a price demand greater than the buyer's valuation, which results in the price being rejected, means the class of buyers is generally worse off when sellers have sufficiently large but limited, rather than infinite, 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 similar buyers in the past and therefore receiving zero payoff, it may also result in the majority of accepted seller price announcements being closer to the buyers' valuation than happens in the case of perfect information. On the other hand, buyers always prefer the sellers to have perfect information regarding the history of play. This suggests a potential role for consumer advocacy groups to assure that information regarding past transactions remains available, and the sellers themselves may not find it optimal to implement a system of perfect information sharing.

3.4 Continuum vs. finite number of sellers

Up to this point, the paper has indirectly assumed that individual sellers ignore the effect their own price announcements may have on future period payoffs. This may result if sellers are drawn from a continuum of agents, are drawn from a finite set of agents but replaced in the set by identical agents after play, or completely discount future period utility (discount rate $\eta = 0$). However, the results presented in this paper may continue to hold when sellers may 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 ignores future periods. When a seller of type α is selected to play in period t, the agent solves:

$$\max_p u_\alpha(p)\left[1 - F_\alpha(p - \delta \mid \check{p}_t, \hat{p}_t)\right] + B(\eta_\alpha, N, p)$$

where $B(\eta, N, p)$ represents the expected, discounted increase in future periods of the game from announcing price p instead of the price p^o, where p^o is the price that α announces when future period payoff is ignored. Technically, $p^o = \arg\max_p u_\alpha(p)\left[1 - F_\alpha(p - \delta \mid \check{p}_t, \hat{p}_t)\right]$. Given the structure of the game previously described, it follows that $B(\eta, N, p^o) = 0$; and $B(\eta, N, p) > 0$ if and only if $\eta \in (0,1]$ and price demand p results in higher expected future period payoff compared to p^o. Additionally, $B(0, N, p) = 0$; $\frac{\partial B}{\partial \eta} > 0$; $B(\eta, N, p) \to 0$ as $N \to \infty$; and $\frac{\partial B}{\partial N} < 0$.[14]

As $\eta \to 0$ or $N \to \infty$, 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 $p \in (\check{p}, \hat{p})$ will always result in at least as high of expected future utility compared with announcing price $p = \check{p}$. This is because choosing $p \in (\check{p}, \hat{p})$ can result in a smaller potential price range in future periods, decreasing

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

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 $\tilde{p} = P_\beta$, only that $\tilde{p} \in [p^*, P_\beta]$. When $\tilde{p} < P_\beta$, 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 $\{F_\alpha(\cdot), u_\alpha(\cdot)\}_{\alpha \in A}$ such that increased risk seeking, or limited information result in improved seller utility or aggregate class earnings. It follows that allowing for a finite set of sellers who care about the strategic consequences of their price announcements does not change the model's results.

References

[1] Blume, L. and D. Easely, 1992, "Evolution and Market Behavior," *Journal of Economic Theory*, 58, 9–40.

[2] Foster, D. and H.P. Young, 1990, "Stochastic evolutionary game dynamics," *Theoretical Population Biology*, 38, 219–232.

[3] Tirole, J., 1988, *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.

[4] Vega-Redando, F., 2003, *Economics and the Theory of Games*, Cambridge University Press, Cambridge, UK.

[5] Young, H.P., 1993, "An Evolutionary Model of Bargaining," *Journal of Economic Theory*, 59, 145-168.

[6] Young, H.P., 1993b, "The Evolution of Conventions," *Econometrica*, 61n1, 57-84.