Cournot Competition as a Continuous State Space Markov Chain

Alexander Halsey

The University of Texas at Austin

In partial fulfillment of the requirements for graduation with the Dean's Scholars Honors Degree in Mathematics

Under the supervision of Dr. Joseph Neeman and Dr. Maxwell Stinchcombe

May 21, 2021

Abstract

The Cournot competition is a game in which two firms vie to produce the optimal quantity of a good. Perfectly rational and fully informed firms would produce the quantities given by the Nash equilibrium, the point at which neither firm could improve their payoff by changing their action. Although the Nash equilibrium for the Cournot competition is well understood, there are several proposed models describing how firms that are not fully informed or perfectly rational might still learn the Nash equilibrium. Two commonly used models are fictitious play and the successive best response strategy. I build on these by using a Markov chain, a model for the evolution of random systems, to capture the probabilistic behavior of imperfect firms. Most of the theory and applications of Markov chains deals with finite or countable state spaces. In order to make sense in the context of game theory, the theory of Markov chains on arbitrary state spaces must first be presented. I will provide the relevant results for general state space theory, then describe its novel applications for learning Nash equilibria in the Cournot model.

Keywords: Markov chain, general state space, Nash equilibrium, Cournot competition

Contents

1	Bac	ckground 4						
	1.1	Gam	e Theory	4				
	1.2	2 Cournot Competition						
	1.3	Learning Models in Repeated Games						
	1.4	Markov Chains on Discrete State Spaces						
	1.5	Markov Chains on General State Spaces						
2	Pro	posed Learning Models						
	2.1	Impe	rfect Responses	10				
		2.1.1	Defining the Markov Chain	10				
		2.1.2	Irreducibility	12				
		2.1.3	Recurrence and Aperiodicity	14				
		2.1.4	Invariance	15				
		2.1.5	Normally Distributed Errors	18				
		2.1.6	Bounding the Stationary Measure	21				
		2.1.7	Uniformly Distributed Errors	27				
	2.2	Stochastic Fictitious Play						
		2.2.1	Defining the Markov Chain	29				
		2.2.2	Properties of the Chain	30				
		2.2.3	Simulations	31				
3	\mathbf{Dis}	Discussion						
4	Acknowledgements							
R	leferences							

1 Background

1.1 Game Theory

Game theory is the study of strategic interactions between rational decision-makers. A rational decision-maker always decides to maximize their expected utility, given their information. Each rational agent is aware that all the other agents they compete against are expected utility maximizers.

Definition 1.1. A game Γ is a collection $(S_i, u_i)_{i \in I}$ where

- I is a non-empty set of agents
- for each $i \in I$, S_i is i's set of possible strategies
- for each i ∈ I, u_i : S → ℝ is i's utility function, where S := ×_{i∈I}S_i is the product of the actions of all the players

The simplest type of game is a one-shot, simultaneous move game, in which each player i simultaneously chooses an action $s_i \in S_i$ and receives a payoff $u_i(s)$, $s := \times_{i \in I} s_i$, that depends on the strategies of all the players in the game. I will use s_{-i} to denote the strategy profile of all players except player i. Simple games can be represented using a payoff matrix, which specifies the payoffs for each player for every possible strategy profile. An example is the Prisoner's Dilemma, in which two players both have the choice to either cooperate or defect. Their payoffs can be represented by the following matrix:

$$\begin{array}{ccc} C & D \\ C & 2,2 & 0,4 \\ D & 4,0 & 1,1 \end{array}$$

Here the first entry in each cell corresponds to the row player's utility, and the second to the column player's. If the players both play C, they each receive a payoff of 2. If the row player plays D and the column player plays C, then the row player receives a payoff of 4 and the column player receives a payoff of 0.

Definition 1.2. A pure strategy Nash equilibrium for Γ is a vector of strategies, $s^* \in S$ such that for all $s_i \in S_i$, $u_i(s_i^*, s_{-i}^*) \ge u_i(s_i, s_{-i}^*)$. This definition captures the notion of mutual best responses: given the strategies of all the other agents, no agent can benefit by changing strategies. For example, in the Prisoner's Dilemma discussed above, (D, D) is the unique Nash equilibrium. No player can benefit by switching to C if their opponent is playing D, and any player will benefit by switching from C to D no matter what their opponent does. In this case, we say the strategy D**strictly dominates** C. This is the source of the dilemma - (C, C) is a better outcome for everyone compared to (D, D), but no strictly rational actor would play C. Mixed strategy Nash equilibria also exist but will not be discussed in this paper.

1.2 Cournot Competition

Definition 1.3. Cournot competition is a game theoretic model describing quantity competition between two firms, where each firm i decides to produce an amount x_i of a homogeneous good. Market inverse demand is set by some function depending on the sum $x_1 + x_2$.

In this case, letting $x := x_1 + x_2$, I will take the inverse demand function p to be p(x) = A - Bx for A, B > 0. I will also assume that the firms are producing quantities x_1, x_2 at a constant marginal cost C, where C < A, which means their profit functions π are given by

$$\pi_i(x_i, x_j) = x_i(A - B(x_i + x_j)) - Cx_i$$

In the Cournot competition, each firm wants i to choose the right quantity x_i^* to produce to maximize its profit function. In this case, we can easily differentiate the profit functions and solve for the maximum to see that

$$(x_1^*, x_2^*) = \left(\frac{A - C - x_2^*}{2B}, \frac{A - C - x_1^*}{2B}\right)$$
$$x_1^* = \frac{A - C}{2B + 1}$$
$$x_2^* = \frac{A - C}{2B + 1}$$

This vector (x_1^*, x_2^*) is the Nash equilibrium for this problem, as both firms are maximizing their own profits given the quantity produced by the other. The symmetry is due to the assumption that both firms face the same demand function and marginal cost. The latter assumption can be relaxed without complicating the solution. This computation also gives the "best response" correspondence. If firm j produces x_j , then firm i's best response is to produce

$$\frac{A - C - x_j}{2B} \tag{1}$$

1.3 Learning Models in Repeated Games

Although the Nash equilibrium is a well-studied concept, there are several competing models for how agents repeatedly playing the same game might learn what the equilibrium is. In the real world, decisions about how much of a good to produce are not one-shot, simultaneous move games. Rather, they can be more realistically modeled as a sequence of one-shot, simultaneous move games, in which firms decide how much to produce at each time-step given the past history of production. We should expect them to gradually learn what the profit-maximizing quantity to produce is.

One learning model for Cournot competition is the successive best response model. In this model, each firm *i* chooses the optimal response to the quantity produced by firm *j* in the previous time-step. As this process continues, the firms gradually learn the equilibrium quantities. The proposition below requires B > 1/2, which will be assumed for the rest of this paper.

Proposition 1. If B > 1/2, the successive best response model for the Cournot competition specified above converges to the unique Nash equilibrium.

Proof. We can view the Nash equilibrium as the fixed point for the best response correspondence given above. The mapping $x_i^t \mapsto x_i^{t+2}$ can be calculated to be $x_i^{t+2} = k + \frac{x_i^t}{(2B)^2}$, where $k = \frac{(2B-1)(A-C)}{(2B)^2}$. If B > 1/2, this is a contraction mapping, so it converges to a unique fixed point x_i^* by Banach's fixed point theorem. That is, the sequences of quantities produced at times $\{0, 2, 4, \ldots\}$ and at times $\{1, 3, 5, \ldots\}$ both converge to the same fixed point (by uniqueness), so the sequence of quantities produced by firm *i* converges to x_i^* . This fixed point must be the Nash equilibrium quantity for firm *i*. Since both firms use the same adjustment process, the vector of quantities produced converges to (x_1^*, x_2^*) .

Another common learning model is deterministic fictitious play. Suppose firm i produces

quantities (x_i^1, \ldots, x_i^T) at times 1 through T. For all times t > T, $x_i^t = Br_i(\frac{1}{t}\sum_{n=1}^t x_j^n)$, where the best response function $Br_i(x_j) = \frac{A-C-x_j}{2B}$ was given earlier. This model corresponds to a pattern where firms compete until time T, at which point they start best responding to the average of each other's past behavior.

Proposition 2. If the best response function is a contraction mapping, deterministic fictitious play converges to the unique Nash equilibrium.

Proof. See Deschamps [2].

These models assume that firms are fully aware of each other's behavior and that they always respond perfectly to one another. These assumptions are unrealistic for a model that attempts to capture behavior in the real world, and I will revisit them later.

1.4 Markov Chains on Discrete State Spaces

A Markov chain is a model for the evolution of a random process. The model is composed of the set of states the system or process can be in at any time, and a set of probabilities governing the transitions between states over time.

Definition 1.4. A discrete time Markov chain is a sequence $\Phi = \{\Phi_n, n \in \mathbb{Z}^+\}$, where Φ_i are random variables on a state space X with the property that given Φ_n, Φ_{n+1} is independent of $\Phi_0, \ldots, \Phi_{n-1}$.

In the simplest case, the state space X contains only a finite or countable (i.e. discrete) number of states. The next state of the chain depends only on the current state of the chain, not its previous history. The values, or sample paths, of the chain Φ are points in the product space $\Omega = \prod_{i=1}^{\infty} X_i$, where each X_i is an exact copy of X. If we let μ be the initial distribution of the Markov chain, and $P_{x_0}(\Phi_1 = x_1)$ be the probability of transitioning from the state x_0 to the state x_1 , then the 'forgetfulness' property of the chain can be formally stated as follows. For every n, and for every sequence of states $\{x_0, \ldots, x_n\}$,

$$P_{\mu}(\Phi_0 = x_0, \Phi_1 = x_1, \dots, \Phi_n = x_n) = \mu(x_0)P_{x_0}(\Phi_1 = x_1)P_{x_1}(\Phi_1 = x_2)\dots P_{x_{n-1}}(\Phi_1 = x_n)$$

 Φ is called a **time-homogeneous** chain if the transition probabilities $P_{x_i}(\Phi_1 = x_{i+1})$ are independent of the times i, i + 1. I will only consider discrete time, time homogeneous Markov chains. For a time-homogeneous Markov chain, $P(x, y) := P_x(\Phi_1 = y)$. Note that the forgetfulness property of the Markov chain can also be stated as

$$P_{\mu}(\Phi_{n+1} = x_{n+1} | \Phi_n = x_n, \dots, \Phi_1 = x_1) = P(x_n, x_{n+1})$$

Definition 1.5. The matrix $P = \{P(x, y), x, y \in X\}$ is known as the **transition matrix** for a Markov chain, and satisfies the following for all x, y:

- $P(x,y) \ge 0$
- $\sum_{z \in X} P(x, z) = 1$

Definition 1.6. Given a transition matrix P, the *n*-step transition matrix P^n gives the distribution of states in the chain after n steps. For $A \subseteq X$,

$$P^n(x,A) := \sum_{y \in A} P^n(x,y)$$

We often want to know whether it is possible for the chain to go from one state to another in a finite number of steps. To be precise, we say that if there is some n such that $P^n(x, y) > 0$, then $x \to y$. If $y \to x$ as well, then $x \leftrightarrow y$, or x and y communicate.

Proposition 3. Communication is an equivalence relation, so the equivalence classes $C(x) = \{y : x \leftrightarrow y\}$ cover the state space X, with $x \in C(x)$

Definition 1.7. If C(x) = X for some state x, then the Markov chain is *irreducible*.

Irreducibility means that it is possible to go from any state in the chain to any other.

1.5 Markov Chains on General State Spaces

It is possible to construct Markov chains on continuous state spaces, and indeed much of the theory for discrete space chains generalizes nicely. I will define a Markov chain on a continuous state space following the convention of Meyn and Tweedie [5], building up from the transition probabilities. We will require the state space X to be equipped with a countably generated σ -field, usually the Borel σ -field $\mathcal{B}(X)$. Note the similarity of the definitions and theorems below to those for discrete state space chains.

Definition 1.8. If $P = \{P(x, A), x \in X, A \in \mathcal{B}(X)\}$ satisfies

- for each $A \in \mathcal{B}(X), P(\cdot, A)$ is a non-negative measurable function on X
- for each $x \in X, P(x, \cdot)$ is a probability measure on X

then P is known as a transition probability kernel.

Theorem 1. For any initial measure μ on $\mathcal{B}(X)$ and a transition probability kernel P we can define a Markov chain $\Phi = \{\Phi_1, \Phi_2, \ldots\}$ on $\Omega = \prod_{i=1}^{\infty} X_i$ that is measurable with respect to $\mathcal{F} = \bigvee_{i=0}^{\infty} \mathcal{B}(X_i)$, and a probability measure P_{μ} on \mathcal{F} such that $P_{\mu}(B)$ is the probability of the event $\Phi \in B$ for $B \in \mathcal{F}$; and for measurable $A_i \subseteq X_i$, $i = 0, 1, \ldots, n$ and any n,

$$P_{\mu}(\Phi_0 \in A_0, \Phi_1 \in A_1, \dots, \Phi_n \in A_n) = \int_{y_0 \in A_0} \cdots \int_{y_{n-1} \in A_{n-1}} \mu(dy_0) P(y_0, dy_1) \cdots P(y_{n-1}, A_n)$$

Proof. See Meyn and Tweedie [5].

Irreducibility for Markov chains on discrete/countable state spaces deals with the hitting times of points in the space. Although sets can be reached with positive probability by a general state space Markov chain, we cannot say the same for individual points in the space, so we need a slightly different notion of irreducibility. Fortunately, we can use the analagous notion of ψ -irreducibility for a general state space chain. First, we define the relevant terms. Let Φ be a Markov chain evolving on a general state space X.

Definition 1.9. For any set $A \in \mathcal{B}(X)$ the first return time of A, denoted by τ_A , is

$$\min\{n \ge 1 : \Phi_n \in A\}$$

Definition 1.10. For any state $x \in X$ and set $A \in \mathcal{B}(X)$, the return time probability, written as L(x, A), is $\mathbb{P}_x(\tau_A < \infty)$.

L(x, A) gives the probability that starting from the state x, the chain reaches the set A in a finite number of steps. With these definitions in mind, we use the framework for irreducibility described by Meyn and Tweedie [5].

Definition 1.11. A Markov chain Φ is called ϕ -irreducible if there exists a measure ϕ on $\mathcal{B}(X)$ such that whenever $\phi(A) > 0$, we have L(x, A) > 0 for every $x \in X$.

This definition captures the idea that Φ hits "large" sets with positive probability, regardless of where it starts in the space. ϕ -irreducibility is a relatively weak property; it does not guarantee that sets of ϕ -measure 0 are avoided with probability 1. Thus, we seek a stronger extension of ϕ , given by the following theorem. **Theorem 2.** If a Markov chain Φ is ϕ -irreducible, then there exists a probability measure ψ on $\mathcal{B}(X)$ such that

- 1. Φ is ψ -irreducible
- 2. for any other measure ϕ' , Φ is ϕ' -irreducible iff $\psi \succ \phi'$
- 3. if $\psi(A) = 0$, then $\psi\{y : L(y, A) > 0\} = 0$

The important features of ψ -irreducibility are that ψ can be taken as a probability measure, and that "negligible" sets as measured by ψ are avoided with probability 1. ψ is sometimes known as a **maximal irreducibility measure**.

2 Proposed Learning Models

2.1 Imperfect Responses

2.1.1 Defining the Markov Chain

Both the fictitious play model and the successive best response strategy assume that the firms respond perfectly to one another at each timestep. Suppose instead that the firms are playing a successive best response strategy, where they choose the best response with probability $1 - \epsilon$, where $\epsilon > 0$. With probability ϵ , they choose the "wrong" response. More precisely, at each time t, with probability ϵ firm i decides what quantity to produce by sampling from some distribution ν , where ν is either fixed or dependent only on the quantities produced by both firms in the previous timestep. This model captures the idea of firms acting either irrationally or being subject to shocks outside of their control. They make the rational decision most of the time, but every once in a while they make a random choice. I will assume that firms do not produce negative quantities, and that they will not produce more than $R := \frac{A-C}{B}$ as this would guarantee they lose money given the demand function specified above. Thus, ν is always supported on the compact set [0, R]. There are several reasonable choices one can make for ν . It can be uniform on [0, R], or it can be normally distributed (truncated to [0, R]) such that at time t, $\mathbb{E}[\nu_i] = Br_i(x_j^{t-1})$. The specific choice of ν will matter when simulating behavior. The theoretical results described

below require that ν be a strictly positive measure, meaning that ν assigns non-zero measure to every non-empty open subset of [0, R].

Definition 2.1. The imperfect response transition kernel P is defined as follows. Let $\epsilon > 0$, and $A = (A_1 \times A_2) \in \mathcal{B}([0, R]^2)$. Let the state at time t be $\alpha = (x_1, x_2)$, and let ν_{α_i} be a positive probability measure on [0, R] assigned to $i \in \{1, 2\}$ depending only on the current state α . Then,

$$P((x_{1}, x_{2}), A) = \begin{cases} (1 - \epsilon + \epsilon \nu_{\alpha_{1}}(A_{1}))(1 - \epsilon + \epsilon \nu_{\alpha_{2}}(A_{2})), & \text{if } Br_{1}(x_{2}) \in A_{1}, Br_{2}(x_{1}) \in A_{2} \\ (1 - \epsilon + \epsilon \nu_{\alpha_{1}}(A_{1}))\epsilon \nu_{\alpha_{2}}(A_{2}), & \text{if } Br_{1}(x_{2}) \in A_{1}, Br_{2}(x_{1}) \notin A_{2} \\ \epsilon \nu_{\alpha_{1}}(A_{1})(1 - \epsilon + \epsilon \nu_{\alpha_{2}}(A_{2})), & \text{if } Br_{1}(x_{2}) \notin A_{1}, Br_{2}(x_{1}) \in A_{2} \\ \epsilon \nu_{\alpha_{1}}(A_{1})\epsilon \nu_{\alpha_{2}}(A_{2}), & \text{if } Br_{1}(x_{2}) \notin A_{1}, Br_{2}(x_{1}) \notin A_{2} \end{cases}$$

$$(2)$$

where $(x_i, x_j) \mapsto (Br_i(x_j), Br_j(x_i))$ is the best response correspondence for the Cournot competition described above.

This transition kernel makes precise the "imperfect" behavior described directly above. By construction, the transition kernel is Markovian, since ν_{α_i} and the best response correspondence depend only on the state α .

Theorem 3. For any initial distribution μ , the imperfect response transition kernel P defines a Markov chain on $[0, R]^2$.

Proof. First we must check that P meets the criteria for definition 1.8. The first is immediate. To show the second, fix $\alpha = (x_1, x_2) \in [0, R]^2$. Then,

$$P(\alpha, \cdot) = (1 - \epsilon)^2 + (1 - \epsilon)\epsilon \int d\nu_{\alpha_1} + (1 - \epsilon)\epsilon \int d\nu_{\alpha_2} + \epsilon^2 \iint d\nu_{\alpha_1} d\nu_{\alpha_2} = (1 - \epsilon)^2 + 2(1 - \epsilon)(\epsilon) + \epsilon^2 = 1$$

since ν_{α_i} is a probability measure. By Theorem 1, the transition kernel defines a Markov chain with the given probability law.

Having shown the lemma, we can view the probabilistic Cournot adjustment process between two firms as a Markov chain evolving on the state space $[0, R]^2$. The state at time t is the vector (x_1^t, x_2^t) of quantities produced by firms 1 and 2. The imperfect response transition kernel gives a distribution over possible quantities produced at time t + 1. Although this is a Markov chain for any initial distribution μ , we will take $\mu = \delta_{\alpha}$ for some state α , since real firms produce a fixed quantity, not a distribution over quantities.

I showed above that under the assumption that B > 1/2, the successive best response strategy converges to the Nash equilibrium. The process just defined cannot remain at a single point, because even if the firms arrive at the Nash equilibrium, with probability ϵ they will deviate, which means they will almost surely move away from the Nash equilibrium at some time after arriving to it. However, the theory of general state space Markov chains will allow us to reach equally desirable conclusions about the imperfect response process.

2.1.2 Irreducibility

To show the imperfect response chain defined above is ψ -irreducible, we use T-chains. Although this may seem like more advanced a technique than is necessary, using T-chains allows us to take advantage of the numerous results presented by Meyn and Tweedie [5].

Definition 2.2. Let $a = \{a(n)\}$ be distribution on \mathbb{Z}_+ , and let Φ be a Markov chain with transition kernel P. The $\mathbf{K}_{\mathbf{a}}$ chain with sampling distribution \mathbf{a} is the Markov chain Φ_a with the transition kernel K_a given by

$$K_a(x,A) := \sum_{n=0}^{\infty} P^n(x,A) a(n) \qquad x \in X, A \in \mathcal{B}(X)$$

The sampled chain can be interpreted as the chain Φ , sampled at time points drawn from a.

Definition 2.3. If a is a sampling distribution, and there exists a substochastic transition kernel T satisfying

$$K_a(x, A) \ge T(x, A)$$
 $x \in X, A \in \mathcal{B}(X)$

where $T(\cdot, A)$ is a lower semi-continuous function for any $A \in \mathcal{B}(X)$, then T is called a continuous component of K_a .

Definition 2.4. If Φ is a Markov chain for which there exists a sampling distribution a such that K_a possesses a continuous component T, with T(x, X) > 0 for all $x \in X$, then Φ is a *T***-chain**.

Theorem 4. If Φ is a T-chain and L(x, O) > 0 for all x and open sets $O \in \mathcal{B}(X)$ then Φ is ψ -irreducible.

With these definitions in mind, we show that the imperfect response chain is ψ -irreducible, for a large class of ν_{α} .

Theorem 5. Let Φ be the imperfect response chain defined in 2.1. For a state $\alpha = (x_1, x_2)$ and a set $A = A_1 \times A_2, A \in \mathcal{B}(X)$, define the transition kernel T as follows:

$$T((x_1, x_2), A) = \epsilon^2 \nu_{\alpha_1}(A_1) \nu_{\alpha_2}(A_2)$$

Assume that all the ν_{α} are positive measures chosen so that $T(\cdot, A)$ is lower semi-continuous for all $x \in X$. Then, Φ is a ψ -irreducible T-chain.

Proof. Let $a = \delta_1$ be the sampling distribution used to define the K_a chain. This is a trivial choice - K_{δ_1} has the same transition kernel as Φ . Next, note that for any state α ,

$$T(\alpha, \cdot) = \epsilon^2 \iint d\nu_{\alpha_1} d\nu_{\alpha_2} = \epsilon^2$$

This computation along with the assumptions made in the statement of the lemma show that T is a continuous component of K_{δ_1} . Clearly, since $K_{\delta_1}(x, A) = P(x, A)$, where P is the imperfect response kernel, we conclude that $K_{\delta_1}(x, A) \ge T(x, A)$ for all $x \in X, A \in \mathcal{B}(X)$. Thus, the imperfect response chain is a T-chain. Finally, let $O = O_1 \times O_2$ be an open set in $[0, R]^2$. Let α be any state in the chain. By construction, $P(\alpha, O) \ge \epsilon \nu_{\alpha_1}(O_1) * \epsilon \nu_{\alpha_2}(O_2) > 0$. Then, $L(\alpha, O) \ge \mathbb{P}_{\alpha}(\tau_O = 1) = P(\alpha, O) > 0$. Thus, for all open sets O and states α , $L(\alpha, O) > 0$, so by Theorem 4, the imperfect response chain is ψ -irreducible.

Proposition 4. For the imperfect response chain, the irreducibility measure ψ assigns non-zero measure to every Borel set of non-zero Lebesgue measure.

Proof. Suppose there exists some set $A = A_1 \times A_2 \in \mathcal{B}([0, R]^2)$ of non-zero Lebesgue measure such that $\psi(A) = 0$. Note that $\psi\{y : L(y, A) > 0\} = 0$, by Theorem 2. However, for all states $\alpha \in [0, R]^2$, $P(\alpha, A) \ge \epsilon \nu_{\alpha_1}(A_1) \ast \epsilon \nu_{\alpha_2}(A_2) > 0$. Then, $\{y : L(y, A) > 0\} =$ $[0, R]^2$, which would mean ψ is trivial, a contradiction.

2.1.3 Recurrence and Aperiodicity

 ψ -irreducibility, while useful, is quite an abstract property in the context of a Cournot competition. As argued before, we cannot guarantee this chain will remain at the Nash equilibrium forever. However, ψ -irreducibility allows us to show that the chain will hit any neighborhood of the Nash equilibrium infinitely many times. This seems like a powerful result, but ψ -irreducibility implies the chain will hit *every* open set in the space infinitely many times. Instead, we will focus on hitting times of open neighborhoods containing the Nash equilibrium. Ideally, the chain will get close to the Nash equilibrium quickly and keep returning there often relative to other states.

Definition 2.5. For a Markov chain Φ and a set $A \in \mathcal{B}(X)$, the occupation time is

$$\eta_A := \sum_{n=1}^{\infty} \mathbb{1}_{\{\Phi_n \in A\}}$$

In this instance, we care about whether or not the imperfect response chain hits a given set A infinitely many times, or equivalently that $\eta_A = \infty$.

Definition 2.6. A set A is called **Harris recurrent** if $Q(x, A) := \mathbb{P}_x(\eta_A = \infty) = 1$ for all $x \in A$. A chain Φ is Harris recurrent if it is ψ -irreducible and every set $A \in \mathcal{B}(X)$ of positive ψ measure is Harris recurrent.

To show the imperfect response chain is Harris recurrent, we use the following result without proof.

Theorem 6. A ψ -irreducible T-chain is Harris recurrent if and only if $P_x\{\Phi \to \infty\} = 0$ for every $x \in X$, where $\{\Phi \to \infty\}$ means the chain Φ visits every compact set only finitely often.

Since the imperfect response chain evolves only on a compact set, by the theorem above it is Harris recurrent. As I showed earlier, every set $A \in \mathcal{B}([0, R]^2)$ with positive Lebesgue measure has positive ψ measure, so if the chain starts at $\alpha \in A$, it returns to A infinitely many times with probability one. This, combined with ψ -irreducibility, should imply the chain visits every open neighborhood containing the Nash equilibrium infinitely many times almost surely, regardless of where it starts. **Theorem 7.** If a Markov chain is Harris recurrent, Q(x, B) = 1 for every $x \in X$ and $B \in \mathcal{B}(X)$ such that $\psi(B) > 0$.

Every neighborhood of the Nash equilibrium has positive ψ -measure, so by the theorem above, the imperfect response chain visits every neighborhood of the equilibrium infinitely many times almost surely. Of course, the theorem also implies that the chain visits every open neighborhood infinitely many times, so we look for stronger results using invariance.

2.1.4 Invariance

As a Markov chain $\Phi = {\Phi_n}$ evolves in time, we may expect it to approach a stable configuration, where stability means that for large n, the distribution of Φ_n is constant independently of n.

Definition 2.7. A stationary measure π for a Markov chain with transition kernel P evolving on a state space X is a σ -finite measure π on $\mathcal{B}(X)$ with the property that

$$\pi(A) = \int_X \pi(dx) P(x, A), \qquad A \in \mathcal{B}(X)$$

The interpretation is that for a Markov chain Φ with initial distribution π , each Φ_n will be distributed according to π . Because the imperfect response chain is ψ -irreducible and recurrent, the following result applies.

Theorem 8. Suppose Φ is recurrent and ψ -irreducible. Then, there exists a unique (up to constant multiples) invariant measure π on $\mathcal{B}(X)$ such that π is equivalent to ψ and satisfies for any $A, B \in \mathcal{B}(X)$ where $\psi(A) > 0$,

$$\pi(B) = \int_A \pi(dy) \mathbb{E}_y[\sum_{k=1}^{\tau_A} \mathbb{1}_{\{\Phi_k \in B\}}]$$

This means that the stationary measure of B is proportional to the amount of time spent in B between successive visits of A, assuming the chain starts in A. This theorem guarantees existence of π , but it does not guarantee finiteness. Fortunately, because Φ is a T-chain evolving on a compact set, we can use the criterion of smallness.

Definition 2.8. A set $C \in \mathcal{B}(X)$ is small if there exists an m > 0, a constant $\delta \ge 0$ and a probability measure μ_m such that for all $x \in C$, $B \in \mathcal{B}(X)$,

$$P^m(x,B) \ge \delta\mu_m(B) \tag{3}$$

A central result of the general state space theory is that for a ψ -irreducible chain, every Borel set of positive ψ measure contains a small set. Another important result is that for a T-chain, compact sets are small. Because the imperfect response chain evolves on a compact set, we can use the following result to conclude that the imperfect response chain has a finite stationary measure.

Theorem 9. The invariant measure π is finite if there exists a small set C such that

$$\sup_{x\in C} \mathbb{E}_x[\tau_C] < \infty$$

If we take $C = [0, R]^2$, we immediately conclude that the unique stationary measure π is finite for the imperfect response chain. Since π is equivalent to ψ , and ψ is positive on all Borel sets on non-zero Lebesgue measure, π assigns positive probability to all Borel sets of non-zero Lebesgue measure. In particular, any neighborhood of the Nash equilibrium has positive π measure. Before studying the long term dynamics of the chain, we define one more relevant term. In studying the long term behavior of the chain, we want to avoid scenarios where it cycles through states in some predictable fashion.

Definition 2.9. For a ψ -irreducible Markov chain, the largest d such that there exist disjoint sets $D_1 \dots D_d \in \mathcal{B}(X)$ such that

- for $x \in D_i$, $P(x, D_{i+1}) = 1$, $i = 0 \dots d 1 \mod d$
- the set $N = [\bigcup_{i=1}^{d} D_i]^c$ is ψ -null

is known as the **period** of the chain. If d = 1, the chain is **aperiodic**.

Proposition 5. The imperfect response chain is aperiodic.

Proof. To see why, recall that all the ν_{α} defining the transition kernel P are positive. If there were some cyclic decomposition of Φ into more than one set, then there would be a state α , where $\alpha \in D$, $D = D_1 \times D_2, D \in \mathcal{B}([0, R]^2)$ such that $P(\alpha, D) = 0$. But, $P(\alpha, D) \ge \epsilon^2 \nu_{\alpha_1}(D_1)\nu_{\alpha_2}(D_2) > 0$. Thus, the chain is aperiodic.

Having shown aperiodicity, Harris recurrence, and the existence of a finite stationary measure, we use the following result. **Theorem 10.** Suppose Φ is an aperiodic, Harris recurrent chain, with invariant measure π . If π is finite, then for every initial state $x \in X$,

$$\sup_{A \in \mathcal{B}(X)} |P^n(x, A) - \pi(A)| \to 0$$

For the imperfect response chain, Theorem 10 implies that no matter where we start in the state space, Φ_n for large n is distributed according to π . Intuitively, the chain spends most of its time in sets of relatively large π -measure. If the chain spends most of its time in sets containing the Nash equilibrium, then in some sense the competing firms have still "learned" the optimal quantities to produce, despite their imperfect response patterns. π itself is difficult to calculate, and it should be evident that it will depend on the specific choice of ν_{α} . In the next section, we provide simulations for different choices of ν_{α} . We can also get strong results about the speed of convergence to the stationary distribution. First, we define the relevant metric on probability distributions.

Definition 2.10. The total variation distance between two probability distributions μ and ν on a space X is

$$\|\nu(\cdot) - \mu(\cdot)\| = \sup_{A \in \mathcal{B}(X)} |\nu(A) - \mu(A)|$$

With this definition in mind, we can apply the following result.

Theorem 11. Consider a Markov chain with stationary distribution π . Suppose the minorization condition (3) holds for some $n_0 \in \mathbb{N}$, $\delta > 0$, and probability measure μ , in the special case C = X, i.e. the entire state space is small. Then,

$$\|P^n(x,\cdot) - \pi(\cdot)\| \le (1-\delta)^{\lfloor n/n_0 \rfloor} \tag{4}$$

where |r| is the greatest integer less than or equal to r.

See Roberts [6] for a proof. Since the entire state space is small for the imperfect response chain, convergence (in the sense of total variation distance) to the stationary distribution π is both uniform over all starting states in the chain and relatively fast for all starting states.

2.1.5 Normally Distributed Errors

Assume the ν_{α} for each firm are normally distributed about the best response to their competitor's last action. If the state at time t is $\alpha = (x_1, x_2)$, then $\nu_{\alpha_1} \sim N(Br_1(x_2), \sigma)$ and $\nu_{\alpha_2} \sim N(Br_2(x_1), \sigma)$. Under this model, even if firms make the "wrong" choice with probability ϵ , that choice is unlikely to be far from the optimal one. There is a small issue with the supports, as Gaussian distributions are not supported on compact sets. We can correct this by conditioning on whether or not each draw from the ν_{α} is in $[0, R]^2$. In principle, this conditioning slightly alters the mean. In practice, if we take σ small enough so that for the Nash equilibrium $(x_1^*, x_2^*), |x_i^* \pm 3\sigma| \in [0, R]$ for each *i*, then the adjustment to the mean is negligible.

We need to check two properties before applying the results from the previous section. First, $\nu_{\alpha_i} = N(Br_i(x_j), \sigma)$ is strictly positive for all states $\alpha = (x_1, x_2)$, so the irreducibility results apply. Second, to show the imperfect response chain is a T-chain, we need to confirm that for all sets $A = A_1 \times A_2$, $T(\cdot, A) = \epsilon^2 \nu_1(A_1)\nu_2(A_2)$ is a lower semi-continuous function (see Proposition 5, and note that ν_1 and ν_2 depend on the state (\cdot)). This is immediate, because as $x_n \to x$, the best response correspondence $Br_i(x_n) \to Br_i(x)$, which means that the distributions $\nu_{\alpha_{i,n}}$, normally distributed about the best responses, converge on the measurable sets A_1, A_2 . Thus, for this choice of ν_{α} , the imperfect response chain is ψ -irreducible, Harris recurrent, and uniformly ergodic with a finite stationary measure.

We can approximate¹ the stationary distribution by allowing the Markov chain to run until its distribution is very close to stationary, then sampling states of the chain. We provide several examples below. In all cases, the parameters of the Cournot model were chosen as A = 100, B = 2, C = 5, which means the Nash equilibrium quantity x^* is 19.

The simulations show that the stationary measure is always concentrated around the Nash equilibrium quantity, indicating that the firms do "learn" the optimal quantity to produce, though they continue to make mistakes. The effects of ϵ and σ are as expected as well: as both ϵ and σ increase, the response chain spends less time around the Nash equilibrium.

¹Code for all the simulations can be found at https://github.com/alexhalsey/cournot_simulations



Figure 1: Simulated stationary distributions of imperfect response process with error rate $\epsilon = 0.1$, and standard deviations $\sigma = 5$ (left) and $\sigma = 2.5$ (right) for the normal distributions ν_{α} . For $\sigma = 5$, $\pi(B_1(x^*)) = 0.88$, and for $\sigma = 2.5$, $\pi(B_1(x^*)) = 0.92$.



Figure 2: Simulated stationary distributions of imperfect response process with error rate $\epsilon = 0.25$, and standard deviations $\sigma = 5$ (left) and $\sigma = 2.5$ (right) for the normal distributions ν_{α} . For $\sigma = 5$, $\pi(B_1(x^*)) = 0.71$, and for $\sigma = 2.5$, $\pi(B_1(x^*)) = 0.81$.



Figure 3: Simulated stationary distributions of imperfect response process with error rate $\epsilon = 0.50$, and standard deviations $\sigma = 5$ (left) and $\sigma = 2.5$ (right) for the normal distributions ν_{α} . For $\sigma = 5$, $\pi(B_1(x^*)) = 0.47$, and for $\sigma = 2.5$, $\pi(B_1(x^*)) = 0.63$.



Figure 4: Simulated stationary distribution of imperfect response process with error rate $\epsilon = 1$, and standard deviations $\sigma = 5$ (top) and $\sigma = 2.5$ (bottom) for the normal distributions ν_{α} . This reflects a scenario when the firms never pick the successive best response, and instead always draw their response from the normal distribution around the theoretical best response. The orange line represents a Gaussian with mean x^* and standard deviation σ . For $\sigma = 5$, $\pi(B_1(x_i^*)) = 0.15$, and for $\sigma = 2.5$, $\pi(B_1(x_i^*)) = 0.30$.

2.1.6 Bounding the Stationary Measure

We can prove the stationary measure is relatively concentrated around the equilibrium for normally distributed ν_{α} using the Wasserstein metric.

Definition 2.11. The *p*-Wasserstein distance between two probability measures μ and ν on a metric space (M, d) is

$$W_p(\mu,\nu) = \left(\inf_{\gamma \in \Gamma(\mu,\nu)} \int_{M \times M} d(x,y)^p d\gamma(x,y)\right)^{1/p}$$

 $\Gamma(\mu,\nu)$ is the set of **couplings** of μ and ν , i.e. the set of measures on $M \times M$ whose marginal distributions are μ and ν .

This metric is sometimes known as the "earth mover's distance", because if μ and ν are viewed as "piles" of mass, $W_1(\mu, \nu)$ represents the minimum cost required to turn μ into ν . The distance term is used because the cost of moving from point a to point b is d(a, b). The following well-known properties of the Wasserstein metric will be useful.

Remark 1. 1. For point masses δ_x , δ_y , $W_1(\delta_x, \delta_y) = |x - y|$.

- 2. For normal distributions $\nu_1 = N(m_1, \sigma_1)$ and $\nu_2 = N(m_2, \sigma_2)$,
 - $W_1(\nu_1,\nu_2) \ge |m_1 m_2|$
 - $W_2(\nu_1,\nu_2)^2 = |m_1 m_2|^2 + \sigma_1^2 + \sigma_2^2 2\sigma_1\sigma_2$
- 3. For any distributions $\mu, \nu, W_1(\mu, \nu) \leq W_2(\mu, \nu)$.
- W_p(μ, ν) = (inf E[d(X,Y)^p])^{1/p} where the infimum is taken over all joint distributions of the random variables X, Y having marginal distributions μ, ν.

From 2 and 3, we conclude that for normal distributions with the same variance ($\nu_1 = N(m_1, \sigma)$ and $\nu_2 = N(m_2, \sigma)$), $W_1(\nu_1, \nu_2) = |m_1 - m_2|$. W_1 is the only metric that will be used in the rest of this paper, so we will simply refer to it as W.

Lemma 1. Consider the imperfect transition kernel with normally distributed errors P. Let P_1 and P_2 be P restricted to the first and second coordinates. Then, for any states (x_1, x_2) and (y_1, y_2) , $W(\delta_{x_1}P_1, \delta_{y_1}P_1) + W(\delta_{x_2}P_2, \delta_{y_2}P_2) \leq \frac{1}{2b}(W(\delta_{x_1}, \delta_{y_1}) + W(\delta_{x_2}, \delta_{y_2}))$. Proof. Note that $\delta_{x_1}P_1 \sim (1-\epsilon)\delta_{Br_1(x_2)} + \epsilon N(Br_1(x_2),\sigma)$ and $\delta_{x_2}P_2 \sim (1-\epsilon)\delta_{Br_2(x_1)} + \epsilon N(Br_2(x_1),\sigma)$. Then,

$$\begin{split} W(\delta_{x_1}P_1, \delta_{y_1}P_1) &= W((1-\epsilon)\delta_{Br_1(x_2)} + \epsilon N(Br_1(x_2), \sigma), (1-\epsilon)\delta_{Br_1(y_2)} + \epsilon N(Br_1(y_2), \sigma)) \leq \\ (1-\epsilon)W(\delta_{Br_1(x_2)}, \delta_{Br_1(y_2)}) + \epsilon W(N(Br_1(x_2), \sigma), N(Br_1(y_2), \sigma))) = \\ (1-\epsilon)|Br_1(x_2) - Br_1(y_2)| + \epsilon |Br_1(x_2) - Br_1(y_2)| = \\ |Br_1(x_2) - Br_1(y_2)| = \\ |Br_1(x_2) - Br_1(y_2)| = \\ |\frac{A-C-x_2}{2b} - \frac{A-C-y_2}{2b}| = \\ \frac{1}{2b}|x_2 - y_2| \end{split}$$

By the same logic, $W(\delta_{x_2}P_2, \delta_{y_2}P_2) \leq \frac{1}{2b}|x_1 - y_1|$. The conclusion follows by adding the two inequalities.

Lemma 2. For any distributions μ_1, μ_2 and $\nu_1, \nu_2, W(\mu_1P_1, \nu_1P_1) + W(\mu_2P_2, \nu_2P_2) \leq \frac{1}{2b}(W(\mu_1, \nu_1) + W(\mu_2, \nu_2))$. That is, the result from the previous lemma holds for arbitrary distributions over states in the chain, not just point masses. μ_i and ν_i represent different distributions over firm i's choices.

Proof. The result from Lemma 1 generalizes to random measures $\delta_{X_1}, \delta_{X_2}, \delta_{Y_1}, \delta_{Y_1}$ where X_1, X_2, Y_1, Y_2 are random variables. To see why, simply note that $\delta_{X_1}P_1 \sim (1-\epsilon)\delta_{Br_1(X_2)} + \epsilon N(Br_1(X_2), \sigma)$, so the same calculations from the previous lemma are still valid. Thus,

$$W(\delta_{X_1}P_1, \delta_{Y_1}P_1) + W(\delta_{X_2}P_2, \delta_{Y_2}P_2) \le \frac{1}{2b}(W(\delta_{X_1}, \delta_{Y_1}) + W(\delta_{X_2}, \delta_{Y_2})) = \frac{1}{2b}(|X_1 - Y_1| + |X_2 - Y_2|)$$

Now, consider distributions μ_1, μ_2 and ν_1, ν_2 . Recall that $W(\mu, \nu) = \inf \mathbb{E}|X - Y|$ where the infimum is taken over joint distributions of X, Y having marginals μ and ν . Since [0, R] with the standard metric is a Polish space, there exist optimal couplings (X_1, Y_1) and (X_2, Y_2) such that $\mathbb{E}|X_1 - Y_1| = W(\mu_1, \nu_1)$ and $\mathbb{E}|X_2 - Y_2| = W(\mu_2, \nu_2)$ [7]. By what was shown above,

$$W(\delta_{X_1}P_1, \delta_{Y_1}P_1) + W(\delta_{X_2}P_2, \delta_{Y_2}P_2) \le \frac{1}{2b}W(\delta_{X_1}, \delta_{Y_1}) + W(\delta_{X_2}, \delta_{Y_2}) = \frac{1}{2b}(|X_1 - Y_1| + |X_2 - Y_2|)$$

We take the expectation of both sides:

$$\mathbb{E}[W(\delta_{X_1}P_1, \delta_{Y_1}P_1) + W(\delta_{X_2}P_2, \delta_{Y_2}P_2)] \le \frac{1}{2b}\mathbb{E}[|X_1 - Y_1| + |X_2 - Y_2|]$$
$$\mathbb{E}[W(\delta_{X_1}P_1, \delta_{Y_1}P_1) + W(\delta_{X_2}P_2, \delta_{Y_2}P_2)] \le \frac{1}{2b}(W(\mu_1, \nu_1) + W(\mu_2, \nu_2))$$

The 1-Wasserstein metric is jointly convex, so $\mathbb{E}[W(\cdot, \cdot)] \ge W(\mathbb{E}[\cdot], \mathbb{E}[\cdot])$. Thus,

$$W(\mathbb{E}[\delta_{X_1}P_1], \mathbb{E}[\delta_{Y_1}P_1]) + W(\mathbb{E}[\delta_{X_2}P_2], \mathbb{E}[\delta_{Y_2}P_2]) \le$$
$$\mathbb{E}[W(\delta_{X_1}P_1, \delta_{Y_1}P_1) + W(\delta_{X_2}P_2, \delta_{Y_2}P_2)]$$

The left hand side is equal to $W(\mu_1 P_1, \nu_1 P_1) + W(\mu_2 P_2, \nu_2 P_2)$ by definition of the expectation of a random measure, so combining the inequalities we conclude

$$W(\mu_1 P_1, \nu_1 P_1) + W(\mu_2 P_2, \nu_2 P_2) \le \frac{1}{2b} (W(\mu_1, \nu_1) + W(\mu_2, \nu_2))$$

Lemma 3. For all $k \in \mathbb{N}$, and all distributions $\mu_1, \mu_2, \nu_1, \nu_2$,

$$W(\mu_1 P_1^k, \nu_1 P_1^k) + W(\mu_2 P_2^k, \nu_2 P_2^k) \le \frac{2R}{(2b)^k}$$

Proof. The result from the previous lemma implies

$$W(\mu_1 P_1^k, \nu_1 P_1^k) + W(\mu_2 P_2^k, \nu_2 P_2^k) \le \frac{1}{2b} (W(\mu_1 P_1^{k-1}, \nu_1 P_1^{k-1}) + W(\mu_2 P_2^{k-1}, \nu_2 P_2^{k-1})) \le \frac{1}{(2b)^2} (W(\mu_1 P_1^{k-2}, \nu_1 P_1^{k-2}) + W(\mu_2 P_2^{k-2}, \nu_2 P_2^{k-2})) \le \dots \le \frac{1}{(2b)^k} (W(\mu_1, \nu_1) + W(\mu_2, \nu_2))$$

For any distributions μ and ν on [0, R],

$$W(\mu,\nu) = \inf_{\gamma \in \Gamma(\mu,\nu)} \int_{[0,R] \times [0,R]} d(x,y) d\gamma(x,y) \le \inf_{\gamma \in \Gamma(\mu,\nu)} \int_{[0,R] \times [0,R]} R d\gamma(x,y) = R$$

Thus,

$$\frac{1}{(2b)^k} (W(\mu_1, \nu_1) + W(\mu_2, \nu_2)) \le \frac{2R}{(2b)^k}$$

We can exploit the symmetry of the problem in the two coordinates to simplify this inequality. Under the assumption that the best response functions (and therefore the equilibrium quantities) are the same for both firms, P_1 and P_2 and their stationary measures π_1 and π_2 are the same. Take $\mu_1 = \delta_{x^*}$, $\mu_2 = \delta x^*$, $\nu_1 = \pi_1$, and $\nu_2 = \pi_2$. This allows us to compare the imperfect response chain starting from the equilibrium quantities (x^*, x^*) to one starting from the stationary measures (π_1, π_2) . Then,

$$W(\delta_{x^*}P_1^k, \pi_1P_1^k) + W(\delta_{x^*}P_2^k, \pi_2P_2^k) \le \frac{2R}{(2b)^k}$$

The two terms on the right hand side are the same, and $\pi_i P^k = \pi_i$, so

$$W(\delta_{x^*} P_i^k, \pi_i) \le \frac{R}{(2b)^k}$$

We can use the inequality above to provide a lower bound on the mass of the stationary distribution close to the equilibrium.

Theorem 12. If $B_r(x^*)$ is a ball of radius r about the equilibrium, then for all $k \in \mathbb{N}$,

$$\pi_i(B_r(x^*)) \ge (1-\epsilon)^{2k} - \frac{R}{r} \frac{1}{(2b)^k}$$

Proof. Suppose both firms start at the Nash equilibrium. At each timestep, the probability that both firms remain at the equilibrium is $(1 - \epsilon)^2$, since the Nash equilibrium is a fixed point of the best response function. Thus, $P^k((x^*, x^*), (x^*, x^*)) \ge (1 - \epsilon)^{2k}$, or when viewed in one coordinate (still assuming the other firm initially produces the equilibrium quantity), $P_i^k(x^*, x^*) \ge (1 - \epsilon)^{2k}$. On the other hand, $W(\delta_{x^*}P_i^k, \pi_i) \le \frac{R}{(2b)^k}$.

Recall that we can think the 1-Wasserstein distance in terms of optimal transport, where $W(\delta_{x^*}P_i^k, \pi_i)$ is the minimal total cost of turning the first distribution into the second. Consider a ball of radius r about x^* . For any coupling or transport plan γ , $\gamma(A \times B)$ is the amount of mass in A that must get sent to B (see Ambrosio and Gigli [1]). In order to turn $\delta_{x^*}P_i^k$ into π_i , we must transport a mass of at least $(1 - \epsilon)^{2k} - \pi_i(B_r(x^*))$ from the Nash equilibrium to $(B_r(x^*))^c$, which is a distance of at least r. To be precise, for any coupling γ of $\delta_{x^*}P_i^k$ and π_i ,

$$\int_{[0,R]\times[0,R]} |x-y|d\gamma \ge \int_{x^*\times B_r(x^*)^c} |x-y|d\gamma \ge \int_{x^*\times B_r(x^*)^c} rd\gamma = r \cdot \gamma(x^*\times B_r(x^*)^c)$$

Since the above is true for all couplings, it is true for the infimum (the Wasserstein distance). Then, for all k,

$$\frac{R}{(2b)^k} \ge W(\delta_{x^*} P_i^k, \pi_i) \ge r[(1-\epsilon)^{2k} - \pi_i(B_r(x^*))]$$

Rearranging the inequality yields the desired result.

The bound provided is valid for all k, so to get the tightest lower bound we optimize over k, i.e. $\pi_i(B_r(x^*)) \ge \max_{k \in \mathbb{N}} (1-\epsilon)^{2k} - \frac{R}{r} \frac{1}{(2b)^k}$. We provide approximate lower bounds on $\pi_i(B_r(x^*))$ for various choices of ϵ and r, given the values of R and b used for the simulations, and compare to the results of the simulations. Note that the lower bounds do not depend on the standard deviation σ of the ν_{α} .

ϵ	r	$\pi_i(B_r(x^*))$ - lower bound	$\pi_i(B_r(x^*))$ - simulation
0.001	1	0.983	0.9985
0.001	2.5	0.985	0.9991
0.001	5	0.986	0.9997
0.01	1	0.875	0.989
0.01	2.5	0.886	0.994
0.01	5	0.895	0.997
0.05	1	0.552	0.939
0.05	2.5	0.589	0.969
0.05	5	0.626	0.985
0.10	1	0.302	0.877
0.10	2.5	0.356	0.932
0.10	5	0.393	0.967
0.15	1	0.150	0.822
0.15	2.5	0.198	0.902
0.15	5	0.235	0.952
0.20	1	0.061	0.763
0.20	2.5	0.094	0.869
0.20	5	0.131	0.936

Table 1: Behavior of the stationary distribution of the quantities produced by firm i for varying error rates ϵ . The errors are normally distributed with standard deviation $\sigma = 5$. The third column represents a theoretical lower bound on the stationary mass of a ball of radius r about the equilibrium x^* , while the fourth represents the simulated mass of that ball.

These results show that for small ϵ , the stationary distribution for the imperfect response chain places a provably large amount of mass very close to the Nash equilibrium. These results are as expected; the limiting case $\epsilon = 0$ corresponds to the successive best response dynamic, which converges to the Nash equilibrium. As ϵ grows, the theoretical bound drops off from the empirical results. However, in all cases the simulations show a large concentration around the equilibrium, indicating the firms learn to produce close to the optimal quantity in most time periods.

2.1.7 Uniformly Distributed Errors

We can also take the ν_{α} to be uniform on [0, R], which reflects a scenario where production shocks are completely random. Note the ν_{α} are all the same, so for simplicity we will just write γ for the uniform distribution on [0, R].

We need to verify a few properties before using the results from previous sections. The uniform measure is strictly positive, and trivially $T(\cdot, A_1 \times A_2) = \epsilon^2 \gamma(A_1) \gamma(A_2)$ is a lower semicontinous function (since it is constant) for all sets $A = A_1 \times A_2$. Thus, the imperfect response chain is ψ -irreducible, Harris recurrent, and uniformly ergodic with a finite stationary measure.

Similarly to the previous section, we can approximate the stationary measure by allowing the chain to run until its distribution is very close to stationary, then sampling states of the chain. The results are pictured below. Again, the parameters of the Cournot model were chosen as A = 100, B = 2, C = 5, which means the Nash equilibrium quantity x^* is 19.



Figure 5: Simulated stationary distribution π of imperfect response process with error rate $\epsilon = 0.10$ for the uniform distribution γ . $\pi(B_1(x^*)) = 0.80$.



Figure 6: Simulated stationary distribution π of imperfect response process with error rate $\epsilon = 0.25$ for the uniform distribution γ . $\pi(B_1(x^*)) = 0.56$.



Figure 7: Simulated stationary distribution π of imperfect response process with error rate $\epsilon = 0.50$ for the uniform distribution γ . $\pi(B_1(x^*)) = 0.27$.

The bias in the simulations occurs because the mean of the uniform distribution on [0, R] is slightly greater than x^* . All the same theoretical bounds determined in the previous section still apply, since we can prove an analog of Lemma 1 for the uniform distribution.

Lemma 4. Consider the imperfect transition kernel with uniformly distributed errors P. Let P_1 and P_2 be P restricted to the first and second coordinates. Then, for any states (x_1, x_2) and (y_1, y_2) , $W(\delta_{x_1}P_1, \delta_{y_1}P_1) + W(\delta_{x_2}P_2, \delta_{y_2}P_2) \leq \frac{1}{2b}(W(\delta_{x_1}, \delta_{y_1}) + W(\delta_{x_2}, \delta_{y_2}))$.

Proof. Note that $\delta_{x_1}P_1 \sim (1-\epsilon)\delta_{Br_1(x_2)} + \epsilon\gamma$ and $\delta_{x_2}P_2 \sim (1-\epsilon)\delta_{Br_2(x_1)} + \epsilon\gamma$. Then,

$$W(\delta_{x_1}P_1, \delta_{y_1}P_1) = W((1-\epsilon)\delta_{Br_1(x_2)} + \epsilon\gamma, (1-\epsilon)\delta_{Br_1(y_2)} + \epsilon\gamma) \le \delta_{Br_1(y_2)} + \epsilon\gamma$$

$$\begin{split} 1-\epsilon)W(\delta_{Br_1(x_2)}, \delta_{Br_1(y_2)}) + \epsilon W(\gamma, \gamma) &= \\ (1-\epsilon)|Br_1(x_2) - Br_1(y_2)| \leq \\ |Br_1(x_2) - Br_1(y_2)| &= \\ |\frac{A-C-x_2}{2b} - \frac{A-C-y_2}{2b}| &= \\ \frac{1}{2b}|x_2 - y_2| \end{split}$$

By the same logic, $W(\delta_{x_2}P_2, \delta_{y_2}P_2) \leq \frac{1}{2b}|x_1 - y_1|$. The conclusion follows by adding the two inequalities.

Since this same result holds, the results from Lemmas 2 and 3 and Theorem 12 hold, so the theoretical lower bounds in table 1 still apply even if the errors are uniformly distributed.

2.2 Stochastic Fictitious Play

(

2.2.1 Defining the Markov Chain

We can also use a Markov chain approach to modify the fictitious play algorithm, which is less myopic and thus arguably more realistic than the successive best response dynamic. We expect firms to keep track of the past when competing, although they will still make mistakes. We can model this behavior by modifying the imperfect response Markov chain constructed in the previous section. We will prove a set of results for fictitious play similar to those proven for the best response dynamic.

Suppose two firms compete until some time T, then begin best responding to the average of each other's behavior over the last T periods with probability $1 - \epsilon$. With probability ϵ , they choose a random response. To define this as a Markov chain requires some notational work, but the underlying idea and results are similar to the imperfect response process. The state space of the chain is the set of possible T-length production histories, i.e. the set of all vectors $((x_1^1, x_2^1), \ldots, (x_1^T, x_2^T))$ for $(x_1, x_2) \in [0, R]^2$. At each timestep, the "oldest" production quantities in the state drop off and are replaced by the new quantities produced at that timestep. This construction allows us to always keep track of the previous Tquantities produced. Definition 2.12. The fictitious play with imperfect responses transition kernel Fis defined as follows. Let $\epsilon > 0$ and $A = (A_1, A_2) \in \mathcal{B}([0, R]^2)$. Let the state at time t be $\alpha = ((x_1^{t-T+1}, x_2^{t-T+1}), \dots, (x_1^t, x_2^t))$, and let ν_{α_i} be a strictly positive probability measure on [0, R] assigned to $i \in \{1, 2\}$ depending only on the current state α . Using the best response correspondence (1), let $(x_i^t)^* = Br_i(\frac{1}{T}\sum_{n=t-T+1}^t x_j^n)$. Then,

$$F(\alpha, ((x_1^{t-T+2}, x_2^{t-T+2}), \dots, (x_1, x_2^t), A) = \begin{cases} (1 - \epsilon + \epsilon \nu_{\alpha_1}(A_1))(1 - \epsilon + \epsilon \nu_{\alpha_2}(A_2)), & \text{if } (x_1^t)^* \in A_1, (x_2^t)^* \in A_2 \\ (1 - \epsilon + \epsilon \nu_{\alpha_1}(A_1))\epsilon \nu_{\alpha_2}(A_2), & \text{if } (x_1^t)^* \in A_1, (x_2^t)^* \notin A_2 \\ \epsilon \nu_{\alpha_1}(A_1)(1 - \epsilon + \epsilon \nu_{\alpha_2}(A_2)), & \text{if } (x_1^t)^* \notin A_1, (x_2^t)^* \in A_2 \\ \epsilon \nu_{\alpha_1}(A_1)\epsilon \nu_{\alpha_2}(A_2), & \text{if } (x_1^t)^* \notin A_1, (x_2^t)^* \notin A_2 \end{cases}$$
(5)

Note that even though the chain technically evolves on a $T \times 2$ (T pairs of quantities) dimensional state space, the first T-1 pairs at time t+1 are determined by simply shifting the state vector at time t to the left by one, since it would not make sense for the history of the chain to change. Thus, $F(\alpha, ((x_1^{t-T+2}, x_2^{t-T+2}), \ldots, (x_1^t, x_2^t), A) = 0$ if $(x_1^{t-T+2}, x_2^{t-T+2}), \ldots, (x_1^t, x_2^t)$ are not the last T-1 entries of α .

This chain is similar to the imperfect response chain - in fact, the imperfect response chain is the fictitious play chain for T = 1. Most of the same results will apply: it is ψ -irreducible, a T-chain under reasonable assumptions, and uniformly ergodic with a finite stationary measure.

2.2.2 Properties of the Chain

Proposition 6. The fictitious play chain is a T-chain, as long as $\nu_{\alpha_i}(A)$ is a lower semicontinuous function of α_i for all $A \in \mathcal{B}([0, R])$.

Proof. Let $\alpha = ((x_1^{t-T+1}, x_2^{t-T+1}), \dots, (x_1^t, x_2^t))$ be the current state of the chain (i.e., the last T production decisions made by each firm). Consider the sampling distribution $a(n) = \delta_T$, i.e. the K_a chain has transition kernel $K_a(\alpha, A) = F^T(\alpha, A)$. Let $A = (A_1^1 \times A_2^1, \dots, A_1^T \times A_2^T)$ be a set in $\mathcal{B}([0, R]^{2 \times T})$. Define the substochastic transition kernel T_F to be $T_F(\alpha, A) = \epsilon^{2T} \nu_{\alpha_1^1}(A_1^1) \nu_{\alpha_2^1}(A_2^1) \cdots \nu_{\alpha_1^T}(A_1^T) \nu_{\alpha_2^T}(A_2^T)$. Here, $\nu_{\alpha_i^t}$ refers to the error distribution ν_{α} that the firm *i* uses at time *t*. Finally, by construction $K_a(\alpha, A) \geq T(\alpha, A)$. T_F is lower semi-continuous as long as $(\alpha_i)_n \to \alpha_i$ implies $\nu_{(\alpha_i)_n}(A) \to \nu_{\alpha_i}(A)$ for all $A \in \mathcal{B}([0, R])$, which was the same condition for the imperfect response chain. \Box

Proposition 7. The fictitious play chain is a ψ -irreducible Harris chain.

Proof. Fix an open set $O = (O_1^1 \times O_2^1, \dots, O_1^T \times O_2^T)$. For any states α , $L(\alpha, O) \geq F^T(\alpha, O) \geq \epsilon^{2T} \nu_{\alpha_1^1}(O_1^1) \nu_{\alpha_2^1}(O_2^1) \cdots \nu_{\alpha_1^T}(O_1^T) \nu_{\alpha_2^T}(O_2^T) > 0$ since the ν_{α} are strictly positive. By Theorem 4, this is enough to conclude the fictitious play chain is ψ -irreducible. The fictitious play chain is a ψ -irreducible T-chain evolving on a compact set, so by Theorem 6, it is Harris recurrent.

Because the fictitious play chain is Harris and ψ -irreducible, there exists an invariant measure π . By Theorem 9, π can be taken as a probability measure since the chain evolves on a compact set (meaning the entire state space is small). To conclude the chain converges to its stationary distribution, we need to show aperiodicity.

Proposition 8. The fictitious play chain is aperiodic.

Proof. We proceed by contradiction. Suppose we have some decomposition $D_1, \ldots D_d$ as in definition 2.9. Consider $D_1 = (C_1^1 \times C_2^1, \ldots, C_1^T \times C_2^T)$. $F(D_1, D_2) = 1$, so it must be true that $D_2 = (C_1^2 \times C_2^2, \ldots, [0, R]^2)$, since the ν_{α} are assumed to be strictly positive measures. But then, since $D_1 \subseteq D_2^c$, $D_1 \subseteq ((C_1^2)^c \times (C_2^2)^c, \ldots, \emptyset)$. This cannot happen (by construction, no set in the state vector can ever be empty), so we conclude that there is no cyclic decomposition for d > 1.

Thus, by Theorem 10, the fictitious play chain converges in the sense of total variation distance to its stationary distribution. By Theorem 11, this convergence is uniform across all starting states. By the same argument used in Proposition 4, ψ places non-zero measure on each set of non-zero Lebesgue measure, which means π does as well since π and ψ are equivalent (Theorem 8).

2.2.3 Simulations

We take the ν_{α} normally distributed about the best responses, as was done in section 2.1.5. To simulate the stationary measure, we allow the chain to run for a long time and then sample states. The parameters of the Cournot model were again chosen as A = 100, B = 2, C = 5, which means the equilibrium quantity is $x^* = 19$. We provide simulations for a few different choices of the error rate (ϵ) and the amount of time periods considered (T). In all cases, ν_{α} are normally distributed with $\sigma = 5$.



Figure 8: Simulated stationary distributions of fictitious play process with error rate $\epsilon = 0.10$ for T = 5 (left) and T = 20 (right). In both cases, $\pi(B_1(x^*)) = 0.915$.



Figure 9: Simulated stationary distributions of fictitious play process with error rate $\epsilon = 0.25$ for T = 5 (left) and T = 20 (right). For T = 5, $\pi(B_1(x^*)) = 0.786$, and for T = 20, $\pi(B_1(x^*)) = 0.791$.



Figure 10: Simulated stationary distributions of fictitious play process with error rate $\epsilon = 0.50$ for T = 5 (left) and T = 20 (right). For T = 5, $\pi(B_1(x^*)) = 0.571$, and for T = 20, $\pi(B_1(x^*)) = 0.578$.

We also provide a table with simulated values for $\pi(B_r(x^*))$ for various choices of ϵ, r , and T. These results show the stationary measure for the imperfect fictitious play chain is very concentrated about the equilibrium. As expected, the mass about the equilibrium decreases as the error rate ϵ increases. The results indicate a benefit to using fictitious play rather than successive best responses, although there is a sort of plateau effect where averaging over more periods does not necessarily improve the performance of the learning model. This requires further investigation - I believe it might be possible to show that the size of ϵ places an upper bound on how concentrated the stationary measure can be no matter how large the look-back window T is. Nevertheless, this table shows that firms learn to play the Nash equilibrium very often under the imperfect fictitious play learning model.

ϵ	Т	$\pi_i(B_1(x^*))$	$\pi_i(B_{2.5}(x^*))$	$\pi_i(B_5(x^*))$
0.001	1	0.9985	0.9991	0.9997
0.001	5	0.9992	0.9993	0.9997
0.001	20	0.9989	0.9992	0.9995
0.01	1	0.989	0.994	0.997
0.01	5	0.991	0.993	0.997
0.01	20	0.991	0.994	0.997
0.05	1	0.939	0.969	0.985
0.05	5	0.958	0.969	0.984
0.05	20	0.958	0.969	0.984
0.10	1	0.877	0.932	0.967
0.10	5	0.915	0.934	0.968
0.10	20	0.915	0.938	0.968
0.25	1	0.708	0.836	0.919
0.25	5	0.786	0.845	0.919
0.25	20	0.790	0.846	0.920
0.50	1	0.469	0.676	0.838
0.50	2.5	0.571	0.692	0.841
0.50	5	0.578	0.692	0.841

Table 2: Behavior of the stationary distribution of the quantities produced by firm i under fictitious play for varying error rates ϵ and look-back windows T. The errors are normally distributed with standard deviation $\sigma = 5$. Each column represents the simulated stationary mass of a ball of radius r about the equilibrium x^* , for r = 1, r = 2.5, and r = 5.

3 Discussion

We demonstrate that the learning process for a Cournot competition can be described as a Markov chain on a continuous state space. The different Markov chain constructions allowed us to study commonly used learning models while relaxing the assumption of perfect rationality. We show that even under these assumptions, firms engaged in a Cournot competition will still 'learn' the Nash equilibrium, in a slightly weaker sense than what is usually intended. Although their imperfect responses prevent them from learning to always play the equilibrium, the simulations and ergodicity results show that the firms produce quantities very close to the equilibrium at most timesteps, and they learn to do so relatively quickly.

One interesting question is whether or not the Markov chain underlying a realistic learning model is time-stationary. If firms arrive close to the Nash equilibrium and begin to notice that producing the equilibrium quantity leads to higher profits, then they might discard whatever learning model they have been using in favor of always producing a fixed quantity close to or at the equilibrium. On one hand, the simulations show that this is basically what happens with learning models proposed in this paper. In addition, even if firms "know" the equilibrium, they are still prone to making mistakes or suffering production shocks. On the other hand, we could define this learning model differently as having different regimes, where as time goes on the probability of making the correct choice increases.

The Markov chain approach allows us to view this problem in even more generality. We can discard the assumption of fixed marginal cost and instead model the marginal cost each firm faces as a random variable. Since both firms should be vulnerable to the same cost shocks, we might model the marginal costs the two firms face as either the same random variable or a pair of correlated random variables. We could also assume the firms have incomplete memory in fictitious play, where they sample from some subset of their competitor's past actions rather than just seeing the past T quantities. Kaniovsky and Young show strong convergence results for the fictitious play with sampling approach in 2×2 games [4], which may generalize to Cournot competition.

The game theory results in this paper were driven in large part by the work of Fudenburg and Levine [3], who give a thorough survey of fictitious play and present an outline for using Markov chains to study strategy profiles of iterated games. Several authors also show how Markov processes can be used to select between strict equilibria in iterated games. In future, the continuous state space theory presented here could be used to provide a different flavor of selection results for more general game theoretic models with more than one strict equilibrium.

4 Acknowledgements

I would like to thank Dr. Joe Neeman and Dr. Maxwell Stinchcombe for their guidance on this project and their patience in helping me. I am also grateful to my friends and family who have supported me through my time at UT. Finally, I would like to thank the Dean's Scholars Honors program for the opportunity to complete this project and for all the benefits, both personal and academic, it has provided me.

References

- Luigi Ambrosio and Nicola Gigli. A User's Guide to Optimal Transport, pages 1–155.
 Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [2] Robert Deschamps. An algorithm of game theory applied to the duopoly problem. European Economic Review, 6(2):187–194, 1975.
- [3] Drew Fudenberg. The Theory of Learning in Games. MIT Press, Cambridge, MA, 1998.
- [4] Yuri M. Kaniovski and H.Peyton Young. Learning dynamics in games with stochastic perturbations. *Games and Economic Behavior*, 11(2):330–363, 1995.
- [5] Sean Meyn, Richard L. Tweedie, and Peter W. Glynn. Markov Chains and Stochastic Stability. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2009.
- [6] Gareth O. Roberts and Jeffrey S. Rosenthal. General state space markov chains and mcmc algorithms. *Probability Surveys*, 1(none), Jan 2004.
- [7] Feng-Yu Wang. Coupling and applications, 2010.