Dynamic Price Competition and Evolutionary Behavior with Search*

Blake Allison
Emory University

Michael Sacks
Clarkson University

Abstract

This paper studies a model of undirected consumer search with boundedly-rational agents. Consumers observe one price, can engage in costly search to learn the other prices, and then purchase from the firm with the lowest observed price. Each consumer searches only if the observed price exceeds their reservation price, which they dynamically update through one of two revision protocols: myopic best responses or imitation. Firms myopically optimize given the current distribution of reservation prices. Short run pricing is characterized by Edgeworth cycles. Each cycle ends when the firm with the larger installed base relents and monopolizes its residual demand. Convergence to the Diamond paradox occurs only if consumers adapt to search sufficiently infrequently. Convergence to a kinked demand equilibrium at the search cost can also occur, as can convergence to the Bertrand paradox. The Bertrand paradox emerges only temporarily when the search cost becomes arbitrarily small. Otherwise, prices cycle indefinitely.

February 26, 2024

JEL Classification: C73, D21, L11, L13

Keywords: Bertrand paradox, bounded rationality, Diamond paradox, Edgeworth cycles, undirected search

---

* *Allison*: Department of Economics, Emory University, Atlanta, Georgia 30322, baallis@emory.edu.
Sacks: School of Business, Clarkson University, Potsdam, New York 13699, msacks@clarkson.edu
We acknowledge and thank Jerrod Anderson for his contributions to an early version of this paper. We are grateful for comments by seminar participants UC Irvine. Declarations of interest: none.

1 Introduction

Over the past several decades, models of consumer search have made tremendous progress in understanding the role of consumers in the market, broadly examining consumer choice based on the information that is freely known and that which may be obtained via costly search.[1] With a handful of notable exceptions, models of undirected consumer search (wherein consumers must incur costs to learn prices) often make use of at least one of two quixotic assumptions: (*i*) that consumers possess equilibrium knowledge or (*ii*) that consumers possess knowledge of firm processes such as pricing strategies.[2] For example, consumers do not observe the price, but know the probability distribution from which prices are drawn. As has been noted by the authors that have established alternative assumptions within their models, it is understood that these assumptions require consumers to possess far more knowledge than can be reasonably expected.

While previous work has sought to limit the informational burden on consumers, we propose to additionally consider limitations on the cognitive burden.[3] We import tools from evolutionary game theory to model these limitations within a framework of dynamic price competition with consumer search. Using this model, we completely characterize both the short- and long-run pricing patterns of the firms and the search behavior of the consumers. Our model blends Stahl's (1989) static framework of undirected consumer search with Maskin and Tirole's (1988) dynamic framework of sticky prices. We adapt these ideas

[1]The literature on undirected consumer search (Stigler, 1961) and clearinghouses (Salop and Stiglitz, 1977) was born in opposition to common knowledge of pricing, studying firm and consumer behavior in markets where consumers observe only a subset of the firms' prices. Consumers may then engage in costly search, either sequentially or all at once, to learn the remaining prices. See, e.g., Stahl (1989), Fershtman and Fishman (1992), Benabou and Gertner (1993), Dana Jr. (1994), Bikhchandani and Sharma (1996), Anderson and Renault (1999), Baye et al. (2006), Arbatskaya (2008), Yang and Ye (2008), Tappata (2009), Janssen et al. (2011), Cabral and Fishman (2012), Garcia et al. (2017), Armstrong (2017), and Preuss (2023).

[2]Instead of assuming that consumers know the probability distribution of prices from which they're searching, Rothschild (1974) assumes that searchers learn about the distribution while they search it. Benabou and Gertner (1993) studies search market equilibria with Bayesian learning (adaptive search and strategic pricing). Bikhchandani and Sharma (1996) studies the optimal stopping rule when the distribution of prices is unknown to searchers. Rauh (1997) studies search when agents have beliefs based on finitely many moments of the distribution of prices and their past market experiences. Lewis (2011) assumes consumers form expectations of prices based on the observed prices during previous purchases. Janssen et al. (2017) develops a Bayesian framework where consumers form and update beliefs regarding the firms' marginal costs. Additionally, Choi et al. (2018) studies the situation in which prices are known but consumers search to learn about their valuation of the good.

[3]See Ellison (2006) and Spiegler (2014) on the use of bounded rationality in industrial organization.

into a continuous time setting in which both consumers and firms are boundedly rational. While this paper studies the duopoly case, the results can be generalized to any finite number of firms.[4] The traditional approach has consumers form rational expectations about the firms' equilibrium price distributions and then search if the expected savings exceeds the cost of search. The formation of such expectations is generally untenable as consumers often lack the necessary knowledge of the firms' production processes or pricing policies. Instead, we use an approach from evolutionary game theory and assume that consumers make their decisions according to simple rules of thumb and periodically adjust their guiding rule via processes that need not require substantial informational or cognitive burdens (Sandholm, 2010). To offer an illustration of our environment, consider a local retail gasoline market. A consumer driving by a gas station observes a price and, based on two factors – their current price sensitivity and a shock (their current fuel level), may decide to stop and refuel or keep driving to the next gas station to see if the price is lower. Through communication with other drivers over time (or based on their own experiences), consumers update their search behavior to be more or less price sensitive.

The emergent market outcome in our model is varying repeated price wars (Edgeworth cycles), a finding unique to the undirected search literature. Firms sequentially undercut one-another until one firm decides to relent by increasing its price, restarting the price war. Generally, these Edgeworth cycles persist in the long run unless consumers can learn to search sufficiently infrequently. We fully characterize when these cycles persist and when, in the long run, prices converge to well known equilibria including the Diamond paradox (monopoly pricing with no search), a kinked demand equilibrium at the search cost, and the Bertrand paradox (marginal cost pricing).

Edgeworth cycles, which are commonly observed in retail gasoline markets (Castanias and Johnson, 1993; Eckert and West, 2004; Noel, 2007a, 2008; Wang, 2009; Doyle et al., 2010; Zimmerman et al., 2012; Isakower and Wang, 2014), make it difficult for consumers to observe or learn the distribution of prices. With price cycles, prices are generally decoupled from marginal cost, making it less informative for search decisions. The upper bound of

[4] See Online Appendix A.1.

the cycle need not (and generally is not) the monopoly price. The lower bound is generally above both the marginal cost and the search cost, rendering each less informative. Moreover, these bounds shift with aggregate search behavior, making learning from past experiences difficult.

In particular, we assume that consumers engage in search if the observed price exceeds a consumer-specific reservation price, which is a combination of a chosen threshold for acceptance together with an instantaneous, individual-specific random shock.[5] Over time, the consumers adjust their thresholds according to one of two common evolutionary dynamics: imitation of superior strategies or the selection of a myopic best response to the current state.[6]

We assume that firms observe the current distribution of consumers' chosen thresholds, but do not possess knowledge of the process by which the consumers adjust these choices.[7] With this limited knowledge, the firms choose their prices to maximize short term profits.[8] Alternatively, we can think of the firms as facing significant capital constraints and any forward-looking strategy that is distinct from the short-tern strategy must sacrifice current profits for future profits, which the capital constrained firms cannot do. Because the shocks to consumers' reservation prices are individual specific, each consumer searches probabilistically from the perspective of the firms, with the probability of search (weakly) increasing in the observed price. The implied tension between extracting surplus from consumers and inducing consumer search endogenously generates downward sloping demand for each firm's good, independent of the structure of each individual consumer's demand.

In this framework, the consumer dynamic can be simply characterized: consumers search more frequently (choose lower thresholds) when the difference in the firms' prices exceeds the cost of search and search less frequently (choose higher thresholds) otherwise. While by itself this is somewhat obvious, though distinct from adaptive search models such as

[5]See Janssen et al. (2017) for a discussion of consumer search with and without reservation prices.

[6]The qualitative results are unaffected by the choice of revision protocol. Only the rate of convergence changes.

[7]Firms may observe these thresholds via simple surveys. Such surveys would not provide enough information to infer the process by which consumer behavior changes.

[8]This assumption is discussed in more detail in Online Appendix C.1 and we discuss the extension to forward looking firms in Online Appendix C.2.

Lewis (2011), it has important implications for the characterization of the pricing dynamic and its potential for convergence. As noted earlier, the short-run equilibrium dynamic is characterized by Edgeworth cycles, where each firm undercuts its competitor until one firm raises its price to monopolize its residual demand, capturing only those consumers that do not search. The low price firm then raises its price to undercut that firm's relenting price. The process then repeats itself.[9]

Our Edgeworth cycles differ from traditional Edgeworth cycles (Maskin and Tirole, 1988; Wallner, 1999; Eckert, 2003) in a few important ways. First, the residual demand of the higher-priced firm is nonzero because of the probabilistic nature of search, softening the incentive to undercut. Second, when a firm relents after prices have been competed sufficiently low, it does so deterministically and to its own benefit. Again, because there is residual demand for the firm with the higher price, a firm relents because the residual profit exceeds that of undercutting its competitor. This is in contrast to Maskin and Tirole's (1988) model, wherein there is no residual demand, so a firm does not immediately benefit from relenting, leading to a stochastic decision to relent as both firms hope the other will relent first.[10]

Nonzero residual demand provides a greater incentive to relent to the firm with the larger installed base of consumers, that is, the firm whose price is observed by more consumers prior to search. This result provides a testable implication that larger firms will tend to be the first to raise prices during price wars. This outcome is consistent with empirical findings in gasoline markets, e.g., Noel (2007b), Atkinson (2009), and Isakower and Wang (2014). Moreover, we can precisely characterize the point at which a firm will relent by linking the search literature to the literature on capacity-constrained price competition via Gelman and Salop's 1983 judo price, which in this context is defined as the highest price a firm's competitor can set such that the firm would rather monopolize its residual demand instead

[9]Edgeworth cycles were first (informally) predicted by Edgeworth (1925), with the presence of capacity constraints driving the emergence of cycles. The notion was formally examined by Shubik (1959), who found that the equilibrium, while not characterized by cycles, involves price dispersion through mixed strategies. See Vives (1993) for a detailed discussion of the non-existence of pure-strategy equilibrium and indeterminacy of prices in Bertrand-Edgeworth games.

[10]Wallner (1999) finds a deterministic reset driven by cycles being exactly three steps in length, so firms alternate on resets.

of continuing the price war. This offers a unique interpretation of consumer search as a soft capacity constraint in that lowering the price increases sales by less than the amount dictated by the consumers' demand.

Third, the peak and trough of the cycles in our model vary with the consumer search behavior and can thus move independent of costs, as is observed in retail gasoline markets (Noel, 2007a). When consumers have a high propensity to search (have chosen lower thresholds), the prices are competed very low and firms may limit their price when relenting. Alternatively, when consumers have a low propensity to search (have chosen higher thresholds), the firms will relent when the price is still relatively high, as they induce little search by charging a high price.[11] In particular, this means that in our framework, firms do not generally compete the price down to marginal cost.[12] Lastly, the cycles in our model are aperiodic and stochastic. The period length of the cycles is inherently random due to the nature of the price stickiness: opportunities for price revisions are themselves stochastic. The period length is then further affected by the consumers' search behavior as the incentives to undercut versus relent shift.

The potential convergence of the long-run equilibrium dynamic depends on the nature of noise in the consumer dynamic, in particular if it is possible that the aggregate probability of search can converge to (approximately) zero. If consumers can become discouraged from searching, then there is almost sure convergence to monopoly pricing and zero search (the Diamond paradox). If, on the other hand, consumers never become entirely discouraged, i.e., if consumers have a sufficiently high probability of search given the maximum chosen threshold and monopoly pricing, then neither the distribution of consumer thresholds nor the prices set by the firms will converge over time. Instead, the stochastic Edgeworth cycles that characterize the short run will persist indefinitely. As a consequence, marginal cost pricing (the Bertrand paradox) is unstable. If prices were to settle at marginal cost,

[11] In general, the upper bound of the cycles cannot be characterized, as it need not possess a monotonic relationship with the firms' installed base or the distribution of reservation prices, depending on the nature of the noise in the consumer search decisions and structure of individual demand. This is further exacerbated by the fact that residual profit need not be quasiconcave, so there may be multiple residual profit maximizing prices. In special cases, such as uniformly distributed noise and linear demand, the upper bound of the cycle moves with the lower bound. See Online Appendix B.1 and B.2 for details.

[12] Cycling with a trough above marginal cost is also observed in Wallner's (1999) finite horizon model (noted above) and Noel's (2008) model with stochastic marginal costs.

consumers would no longer benefit from searching and thus would gradually adopt higher reservation prices. As the probability of consumer search decreases, one of the firms is able to raise its price and sell a positive quantity since there are some consumers that do not search, thereby obtaining positive profits.

The remainder of the paper is structured as follows. Section 2 presents the model. The comparative statics and dynamics are presented in section 3. Discussion and concluding remarks are provided in section 4.

2 The Model

This section develops a continuous-time duopoly model with search in which firms compete by choosing prices and consumers, endowed with thresholds influencing search decisions, update those thresholds over time. Prices and thresholds are sticky: the firms' and consumers' opportunities to revise their strategies are stochastic and governed by Poisson processes. The primary motivation for utilizing a continuous time approach is that for any open interval of time, the following two mutually exclusive events occur with positive probability: the firms update their prices multiple times before consumers update their thresholds and consumers update their thresholds multiple times before the firms update their prices.[13]

2.1 Preliminaries and Timing

The market consists of two identical firms competing in prices and a continuum consumers with unit mass and identical consumption preferences. A typical firm is indexed by i. Time flows continuously and is indexed by $t \in [0, \infty)$. Denote by $p^t = \left(p_1^t, p_2^t\right)$ the vector of firms' prices at each time t. To avoid confusion, ξ denotes a price that is not associated with any particular firm. The firms have a constant marginal cost of production normalized to zero.

At each time t, market activities occur in four stages.[14]

[13]Though the model is cast in continuous time, it can be interpreted as a sequential model à la Maskin and Tirole (1985) and Maskin and Tirole (1988); however, this requires introducing an extra mechanism for consumer threshold revisions. The Poisson parameters (the rates at which the firms and consumers update) only affect the relative likelihood of each event occurring. Hence, the results are independent of the parameters insofar as the probability of each event occurring is distinguishable from zero and the explicit parameter values are omitted.

[14]The first three stages are akin to the Diamond-Stahl model of undirected search (at each t).

Date 0. Each firm, if given the opportunity to adjust its price via an independent Poisson process, selects its price from a finite grid $G = \{g_0, g_1, \ldots, g_M\}$, where $g_m < g_{m+1}$, to maximize its instantaneous profits.[15] Otherwise, prices are unchanged.

Date 1. Each consumer observes a single firm's price and may search at cost $c > 0$ to learn the other firm's price. The probability of a consumer observing price p_i^t is $\alpha_i \in (0,1)$.[16] Let $\alpha_1 = \alpha \in (0,1)$.

Date 2. After the search decision, each consumer purchases from the firm with the lowest observed price ξ. When both firms set the same price, a searching consumer purchases from the first observed price; i.e., she buys from firm i with probability α_i.

Date 3. Each consumer receives opportunities to update her threshold (detailed below) according to a Poisson process, which is independent across consumers.

The initial price vector p^0 is exogenously fixed, though the equilibrium dynamics do not depend on this starting value. Before proceeding, three remarks are in order. First, the grid G is used to ensure that best responses are well-defined, though throughout we will make comparisons to values defined independently of the grid. Second, the assumption that at each instant, the price observed by each consumer (α_i) is independent of previous purchases does not influence the dynamics or equilibrium outcomes.[17]. Third, the results are not sensitive to the relative positioning of date 3.

2.2 Consumer Preferences and Search Strategies

Each consumer is endowed with the stationary instantaneous utility function $u(q, \xi) = v(q) - \xi q$, where q is the quantity of the good consumed, $v(\cdot)$ is strictly concave, and ξ is the price at which the good is purchased. Under this specification, there exists a continuous decreasing function $D(\xi) = \arg\max_q u(q, \xi)$ that specifies the quantity that each consumer will purchase at the price ξ at any time t. Because the mass of consumers is unity, individual

[15] We discuss this myopic assumption in greater detail in Online Appendix C.1.

[16] The following example helps illustrate why the probability of observing a given price is random. Suppose that there is a city with two gas stations, one located on each side of the city. It is unclear *ex ante* which side of the city a given driver will be on when needing to refuel. The driver observes the price of the closest station.

[17] This assumption simplifies the analysis and notation. We discuss this in greater detail in Online Appendix A.2.

demand $D(\xi)$ corresponds to market demand. It is useful to define the consumers' indirect utility function $v(\xi, s) = u(D(\xi), \xi) - cs$, where $s = 1$ if the consumer searches and $s = 0$ otherwise.

Assumption 1 (A1). *$\xi D(\xi)$ is strictly quasiconcave with unique maximizer ξ^m.*

The consumers' search decision is governed by a simple rule of thumb: a consumer searches if the observed price exceeds some reservation price. Each consumer is endowed with a threshold $\tau \in \{\tau_0, \ldots, \tau_L\}$, where $\tau_k < \tau_{k+1}$ and $L \geq 1$.[18] These thresholds correspond to the consumers' strategies. Define X as the unit simplex in $\mathbb{R}^L$:

$$X := \left\{ (x_0, \ldots, x_L) \in \mathbb{R}^{L+1} : x_k \in [0,1] \text{ for all } k = 0, \ldots, L \text{ and } \sum\nolimits_{k=0}^{L} x_k = 1 \right\}.$$

Denote by $x^t = \left(x_0^t, \ldots, x_L^t\right) \in X$ the mass of consumers endowed with each threshold at time t. Assume $x_k^0 > 0$ for all k, otherwise, any k such that x_k^0 can be removed from the grid and the game proceeds identically.

At each time t, each consumer receives an independent random shock σ^t to her threshold τ and then searches in that period if and only if the observed price ξ exceeds both the perturbed threshold and the search cost; i.e., if $\xi > \max\left\{\tau + \sigma^t, c\right\}$.[19] Let $\tau + \sigma^t$ denote the *minimal acceptable (reservation) price*, where any observed price $\xi > \tau + \sigma^t$ leaves that consumer dissatisfied and willing to search.[20] However, the consumer is still self-serving and recognizes that if the price is less than the cost of search, then the savings from a lower price would not justify the search. The shock σ^t is independent across time and distributed according to the distribution φ. A consumer that observes a price ξ will therefore search with probability $\varphi(\xi - \tau)\iota_{(c,\infty)}$, where $\iota_{(c,\infty)}$ is the indicator function for $\xi \in (c, \infty)$.

Assumption 2 (A2). *$\varphi(\cdot)$ is continuous and strictly increasing on $(-\tau_L, \xi^m - \tau_0)$.*

[18] The results generalize to the case of continuous thresholds following Cheung (2016) and Cheung and Wu (2018).

[19] The assumption that this shock is independent across consumers is unnecessary for the purposes of this paper. It is, however, very plausible and guarantees that the expected profits coincide with actual profits.

[20] The stochastic behavior generated by σ^t is similar to that of the Bayesian model in Janssen et al. (2017), where consumers form and update beliefs regarding production costs.

A2 imposes two relatively mild conditions on the distribution of the shocks to consumers' search thresholds. First, $\operatorname{supp} \varphi \supseteq (-\tau_L, \xi^m - \tau_0)$, though the density remains unrestricted on this interval. Second, the magnitude of the shock may be large enough that a consumer with the highest possible threshold is dissatisfied with any positive price and a consumer with the lowest possible threshold may be satisfied with any price below the monopoly price. This is not a particularly imposing assumption as the probability of these events may be arbitrarily small. Let

$$\bar{\varphi}(\xi, x) = \begin{cases} \sum_{k=0}^{L} \varphi\left(\xi - \tau_k\right) x_k & \text{if } \xi > c \\ 0 & \text{if } \xi \leq c \end{cases}$$

denote the average probability that a random consumer searches after observing a price ξ. Because there are a continuum of consumers, $\bar{\varphi}(\xi, x)$ is equivalently the mass of consumers that search after observing a price ξ. Note that $\bar{\varphi}(\xi, x)$ is strictly increasing in ξ on $[c, \xi^m]$, which consistent with the empirical evidence from gasoline markets (Lewis and Marvel, 2011).[21]

The distribution of the consumers' thresholds x^t evolves as the consumers update their individual strategies. Each consumer receives opportunities to update her threshold according to a Poisson process, which is independent across consumers.[22] We consider two classes of decision rules when consumers change their thresholds, which represent different informational burdens as well as degrees of sophistication in behavior: the best response dynamic and an imitation dynamic. Under the best response dynamic, consumers choose a threshold that is in the set of myopic best responses to the current state (the price vector p^t). This revision protocol does not require explicit knowledge of the prices, only of the payoffs other consumers receive for each possible threshold. Let

$$E\left[v|p^t, \tau\right] = \sum_{i=1}^{2} \alpha_i \left(\varphi\left(p_i^t, \tau\right) v \left(\min_j p_j^t, 1 \right) + \left(1 - \varphi\left(p_i^t, \tau\right)\right) v\left(p_i^t, 0\right) \right).$$

denote the expected utility of a consumer with threshold τ. The evolution of x^t under the best response dynamic is defined by the differential inclusion

$$\dot{x}^t \in B\left(p^t\right) - x^t,$$

[21] We prove this and other supporting results on $\bar{\varphi}(\xi, x)$ in Lemma 1 in the Appendix.

[22] As with the pricing dynamic, the rate of this process is irrelevant and thus omitted.

where $B(p) \subset X$ denotes the best response correspondence for the consumers and $\dot{x}^t$ denotes the derivative of the state with respect to time.

However, this protocol still places a significant informational burden on the consumers. Contrary to the best response dynamic, the imitation dynamic imposes minimal informational burdens on the consumers. Under this dynamic, when a consumer has the opportunity to change her threshold, she is matched uniformly at random with another consumer. Conditional on being matched with a consumer with threshold τ_ℓ, a consumer with threshold τ_k adopts τ_ℓ with probability $r_{k\ell}\left(E\left[v|p^t,\tau_k\right], E\left[v|p^t,\tau_\ell\right]\right)$. The overall probability that a consumer switches from threshold τ_k to τ_ℓ is thus $\rho_{k\ell} = x_\ell r_{k\ell}$.

Assumption 3 (A3). *$r_{k\ell}\left(E\left[v|p^t,\tau_k\right], E\left[v|p^t,\tau_\ell\right]\right) > r_{\ell k}\left(E\left[v|p^t,\tau_\ell\right], E\left[v|p^t,\tau_k\right]\right)$ if and only if $E\left[v|p^t,\tau_k\right] > E\left[v|p^t,\tau_\ell\right]$.*

Given a pair of strategies, A3 requires that consumers are more likely to switch from the strategy that performs worse to one that performs better than the reverse. The evolution of x^t under the imitation dynamic is defined by the system of differential equations

$$\dot{x}_k^t = \sum_{\ell=0}^{L} x_\ell^t \rho_{\ell k} - x_k^t \sum_{\ell=0}^{L} \rho_{k\ell}.$$

for all $k = 0, \dots, L$

2.3 Firm Demand, Pricing, and Profits

We now construct each firm's demand as a function of the prices p and the distribution of consumers' thresholds x. If firm i's price is lower than its competitor's price, then i serves the consumers that initially observe p_i along with all of the searching consumers. If the two firms set the same price, then firm i will serve those consumers that initially observe p_i. Finally, if firm i does not have the lowest price, then it will serve only those consumers that initially observe p_i and do not search. Firm i's demand is thus

$$D_i(p,x) = D(p_i) \times \begin{cases} \alpha_i + (1-\alpha_i)\bar{\varphi}(p_{-i},x) & \text{if } p_i < p_{-i} \\ \alpha_i & \text{if } p_i = p_{-i} \\ \alpha_i\left(1-\bar{\varphi}(p_i,x)\right) & \text{if } p_i > p_{-i}. \end{cases}$$

Let the demand when $p_i < p_{-i}$ be the *front-side demand* and the demand when $p_i > p_{-i}$ be the *residual demand.* The *front-side profit* $\pi_i^F(p, x)$ and *residual profit* $\pi_i^R(\xi, x)$ are defined analogously:

$$\pi_i^F(p, x) = p_i D(p_i)\big(\alpha_i + (1 - \alpha_i)\bar{\varphi}(p_{-i}, x)\big)$$
$$\pi_i^R(\xi, x) = \xi D(\xi)\alpha_i (1 - \bar{\varphi}(\xi, x)) \,.$$

Note that there is positive residual demand facing the firm that does not have the lowest price, even without the presence of capacity constraints. This demand is present due to the stochastic nature by which consumers search. Some consumers will not search and instead purchase at the higher price. Like the consumer search models in Varian (1980) and Stahl (1989), the residual demand is independent of the low price, though in contrast to these models, the front-side demand depends on the high price.

As the marginal cost of production is zero, the revenue function $\xi D(\xi)$ corresponds to the firms' monopoly profit, denoted by $\pi^m = \xi^m D(\xi^m)$, where ξ^m is the monopoly price. For any price $p_{-i} = g_\omega \in G$, denote by $R_i(g_\omega, x)$ firm i's best response correspondence. Observe that R_i constitutes a Markov strategy where the state is given by the competitor's current price g_ω and the current distribution of thresholds x.

3 Equilibrium

As the firm's strategies depend only on the current state and the consumers' revision protocol depends (at most) on the current prices and current state x, our solutions constitute Markov perfect equilibria (MPE).[23] We then explore how x evolves in both the short run and long run under the MPE. We will show below that first order stochastic dominance offers a natural way to characterize the evolution of x. In this context, a distribution x first order stochastically dominates a distribution x' if, for all $k = 0, \ldots, L$,

$$\sum_{\ell=0}^{k} x_\ell \leq \sum_{\ell=0}^{k} x'_\ell,$$

strictly so for at least one k.

[23] In Online Appendix C.2, we discuss how the results extend to the case of forward-looking firms, where Markov perfection is a more substantial condition.

3.1 Residual Maximizers, Judo Prices, and Best Responses

For any distribution of thresholds x, define the set of residual maximizers as $\tilde{P}(x) := \arg\max_\xi \pi_i^R(\xi, x)$ and define the set of residual maximizers on a grid G as $\tilde{P}(x, G) := \arg\max_{g \in G} \pi_i^R(g, x)$. Subscripts on $\tilde{P}(x)$ and $\tilde{P}(x, G)$ are unnecessary as each firm's residual profit function is a constant multiple of the other.

Proposition 1. *Under A1 and A2* $\tilde{P}(x)$ *is nonempty and* $\sup \tilde{P}(x) \leq \xi^m$.

All proofs are contained in the Appendix. With a finite grid of prices, the existence of a residual maximizer in $\tilde{P}(x, G)$ is trivial. However, we will interpret the results as the grid becomes arbitrarily fine, so it is useful to compare the results on the grid to values that are defined independent of the grid.

Given Proposition 1, all prices are henceforth restricted to be weakly below the monopoly price ($p_i^t \leq \xi^m = g_M$) as there is no justification for any firm to price above ξ^m. Even if a firm's price were set above ξ^m, that firm would reduce its price given its first opportunity to do so and would never subsequently increase its price above the monopoly level.

The equilibrium characterization is based on the *judo price* of each firm. A firm i's judo price is the highest price its competitor may set such that i prefers to monopolize its residual demand rather than undercut.[24] Formally,

$$p_i^*(x) := \sup\left\{\xi \leq \xi^m : \xi D(\xi)\left(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x)\right) < \max\nolimits_{p_i} \alpha_i p_i D(p_i)\left(1 - \bar{\varphi}(p_i, x)\right)\right\}, \quad (1)$$

which is defined independently from the grid G. We can similarly describe the judo price when prices are constrained to the grid G:

$$p_i^*(x, G) := \max\Big\{g_\omega \in G \setminus \{g_0\} : g_{\omega-1} D(g_{\omega-1})(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega, x)) \\ \leq \max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x))\Big\}$$

[24] The term judo price originates in a model of entry with sequential pricing developed in Gelman and Salop (1983). The authors draw an analogy between firm strategies and the martial art of judo by pointing out that an entrant firm forces accommodation from the incumbent by setting a low price and limiting its size, thereby incentivizing the incumbent to maintain a large profit margin at a higher price rather than engaging in a price war. Traditionally, a firm's judo price refers to the highest price that it may charge such that the other firm prefers to monopolize residual demand.

In Lemma 2 in the Appendix, we show that $p_i^*(x) \in [c, \xi^m)$. Let $\|G\| := \max_{\omega \geq 1} g_\omega - g_{\omega-1}$ denote the norm of G. We can then characterise the firms' best response correspondences.

Proposition 2. *Suppose A1 and A2 hold. For a sufficiently small but positive δ, if $\|G\| < \delta$, then*

$$R_i(g_\omega, x) = \begin{cases} \{g_{\omega-1}\} & \text{if } g_\omega > p_i^*(x, G) \\ \{g_{\omega-1}\} \cup \tilde{P}(x, G) & \text{if } g_\omega = p_i^*(x, G) \\ \tilde{P}(x, G) & \text{if } g_\omega < p_i^*(x, G). \end{cases}$$

Furthermore, if G^n is such that $\|G^n\| \to 0$ as $n \to \infty$, then

(i) *if $g^n \in \tilde{P}(x, G^n)$ and $g^n \to \xi$ as $n \to \infty$, then $\xi \in \tilde{P}(x)$,*

(ii) *$p_i^*(x, G^n) \to p_i^*(x)$ as $n \to \infty$.*

Proposition 2 demonstrates that when prices are restricted to a grid, each firm's best response correspondence mimics its "best responses" when the space of prices is unrestricted: a firm will undercut it's competitor's price unless that price is below the judo price. This interpretation is only approximate as a best response to prices above the judo price does not exist when the space of prices is a continuum. Note that it is possible that a firm's best response is to set the same price as its competitor. In this case, its competitor's price will be a residual maximizer, and so it will be reflected in the term $\tilde{P}(x, G)$.[25]

As the judo price is a defining feature of the firms' best responses, it will play a large role in the equilibrium dynamics. Thus, it is useful to identify which firm has the lower judo price and how each firm's judo price changes with the distribution of consumer search thresholds.

Proposition 3. *Under A1, if $\alpha_i > \frac{1}{2}$, then $p_i^*(x) \geq p_{-i}^*(x)$. If, in addition to $\alpha_i > \frac{1}{2}$, $p_i^*(x) > c$, then $p_i^*(x) > p_{-i}^*(x)$.*

The firm with the larger installed base is less willing to engage in a price war because having a larger installed base guarantees a greater residual demand and thus higher residual profits. This result is analogous to the result found in studies of capacity constrained price

[25]This statement is formally proven in Lemma 4 in the Appendix.

competition that the firm with the larger capacity has a higher judo price.[26] The following Proposition relates the judo price to the distribution of consumer thresholds.

Proposition 4. *Under A1 and A2, if x first order stochastically dominates x', then $p_i^*(x) \geq p_i^*(x')$. For a sufficiently small and positive δ, if $\|G\| < \delta$, then $p_i^*(x, G) \geq p_i^*(x', G)$.*

As more consumers adopt a higher threshold for search, the judo price increases and as consumers search less frequently, the residual profit increases while the front-side profit decreases, so a firm has less of an incentive to undercut its competitor. While first order stochastic dominance may seem like a restrictive condition to focus on, we will demonstrate shortly that this characterization is the most natural and relevant way to characterize the judo price with respect to consumer behavior.

It merits mentioning that it is generally not possible to demonstrate the same relationship for the residual maximizers. The reason is that the relationship between $\tilde{P}(x)$ and x depends heavily on the functional form of φ. In special cases, the upper bound of the cycles move with the lower bound.[27]

3.2 Equilibrium Dynamics

Consumers will gradually adopt higher search thresholds whenever the prices are such that search has a negative expected profit, and will otherwise gradually adopt lower search thresholds.

Proposition 5. *Suppose that A2 and A3 hold and that $p^t = p$ for all $t \in [T, T + \varepsilon)$ for any time T and $\varepsilon > 0$. Define $c^*(p)$ by*

$$c^*(p) := \sup \{c \geq 0 : v(\min p, 1) > \alpha v(p_1, 0) + (1 - \alpha) v(p_2, 0)\}.$$

For all $t, t' \in [T, T + \varepsilon)$ with $t > t'$, it follows that

(i) if $c < c^*(p)$, then $x^{t'}$ first order stochastically dominates x^t,

[26] See, e.g., Osborne and Pitchik (1986), Deneckere and Kovenock (1992), and Allison and Lepore (2016).

[27] Examples offered in Online Appendix B.1 and B.2 demonstrate that the residual maximizer can move in either direction.

(ii) if $c > c^*(p)$, then x^t first order stochastically dominates $x^{t'}$.[28]

The fact that the distribution of consumer thresholds is always increasing or decreasing (under the ordering induced by first order stochastic dominance) is particularly useful because it implies that Proposition 5 characterizes the motion of the firms' judo prices over time (by Proposition 4).

3.2.1 Edgeworth Cycles

Equilibrium pricing takes the approximate form of cycles of price wars in which firms drive down the price to the point that one firm relents and raises its price, starting the cycle anew. Due to the constantly changing consumer thresholds and stochastic nature of price stickiness, the actual pattern of pricing is not quite cyclical in the classical sense in that the bounds of the price war are not constant. Nonetheless, the general pattern repeats and maintains the characteristics of an Edgeworth cycle. By Proposition 2, the cycles can be formally described using the judo prices as defined on the unconstrained price space. These price cycles will be qualitatively identical on the constrained space for a sufficiently fine grid, the only difference being that the judo prices and residual maximizers may differ by some arbitrarily small amount (bounded by $\|G\|$).

Define $p^*(x) = \max\{p_1^*(x), p_2^*(x)\}$ as the *critical judo price.* The equilibrium prices are described by the following pattern, which resembles an Edgeworth cycle. Consider an initial price vector such that $p_i^0 > p_{-i}^*(x)$ for both firms and suppose that the distribution of consumers' thresholds were to remain fixed. Without loss of generality, assume that firm 2 has a weakly larger initial share of consumers, and thus by Proposition 3 the higher judo price. Hence, firm 2's judo price is also the critical judo price. The Edgeworth cycle proceeds as follows:

1. Firm 1 sets a price just below p_2 and will not adjust it until p_2 changes.
2. Firm 2 sets a price just below p_1 and will not adjust it until p_1 changes.
3. Steps 1 and 2 repeat until the prices are reduced to the critical judo price $p_2^*(x)$.

[28] The case in which $c = c^*(p)$ can be ignored, as the grid may always be perturbed such that no prices satisfy this relationship.

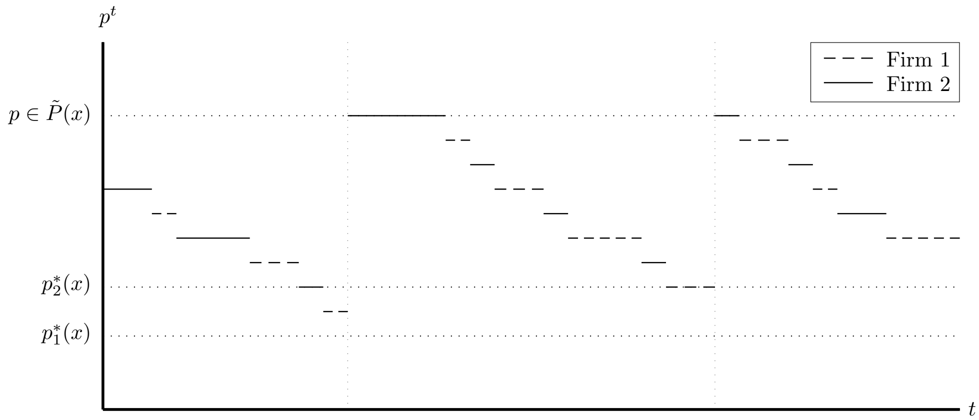

Figure 1: Edgeworth cycles for a fixed distribution of thresholds x.

4. Firm 2 relents and sets a price in $\tilde{P}(x)$ to maximize its residual profit.

5. Repeat this process from step 1.

An example is depicted in Figure 1.

During the first three steps of the cycle, the prices of the firms will be close enough that search will not be beneficial. Thus, Proposition 5 implies that during steps 1-3 of the cycle, consumers will be adopting higher search thresholds, which by Proposition 4 implies that the firms' judo prices will be increasing. If these judo prices increase enough, then a firm that currently has the lowest price (just below the other firm's price) may skip to step 4 of the cycle and monopolize its residual demand. While Proposition 3 guarantees that the firm with the larger initial share of consumers will always relent first for a fixed distribution of consumer thresholds, changes in the distribution can lead to the other firm relenting first. Because the consumer dynamic is continuous, as consumers raise their thresholds and judo prices increase, the critical judo price may increase to a value above the current price before the nonbinding judo price. As such, if the firm with the smaller installed base relents first, then the other firm would also have relented if it had received the opportunity to do so.

The range of equilibrium pricing and potential convergence depend on the parameters and functional forms of the model. This section considers two conditions and the subsequent section analyzes their complements. Denote by

$$e_k = (\underbrace{0, \dots, 0}_{k \text{ zeros}}, 1, 0, \dots, 0).$$

the $k+1^{th}$ basis vector in X. That is, e_k corresponds to the distribution of consumer thresholds in which all consumers have threshold τ_k.

Condition 1 (C1)**.** *If $p_i = \inf \tilde{P}(e_L)$ and $p_{-i} = p_i^*(e_L) = p^*(e_L)$, then $c < c^*(p)$.*

C1 states that the cost of search is sufficiently low so that search is beneficial when all consumers have the highest possible threshold, one firm charges the critical judo price, and the other firm charges the smallest residual maximizer. C1 not only provides an explicit condition on the cost of search, but also implicitly puts some structure on the profit functions in that it requires $p^*(e_L) < \inf \tilde{P}(e_L)$.

The following Theorem demonstrates that under A1-A3 and C1, the equilibrium pricing dynamic does not converge and that the range of prices in the cycles is large enough to induce search. Define $\underline{p} = p^*(e_0)$ and $\hat{p} = \inf \tilde{P}(e_L)$. Thus, $\underline{p}$ represents the critical judo price when all consumers have the lowest search threshold and $\hat{p}$ represents the smallest residual maximizer when all consumers have the highest search threshold.

Theorem 1. *Under A1−A3 and C1, for all $\varepsilon > 0$, there exists a $\delta > 0$ and $T > 0$ such that if $\|G\| < \delta$, then*

(i) neither p^t nor x^t converge as $t \to \infty$,

(ii) for all times $t > T$, $p_i^t \in [\underline{p} - \varepsilon, \xi^m]$,

(iii) $(x^t, p_i^t) \to (e_0, \xi \leq \underline{p} + \varepsilon)$ and $(x^t, p_i^t) \to (e_L, \xi' \geq \hat{p} - \varepsilon)$ infinitely many times.

Thus, the range of prices is at least $[\underline{p}, \hat{p}]$, and the distribution of consumer thresholds varies between the two extremes in which all consumers have the lowest and highest thresholds e_0 and e_L. Proposition 2 implies that some firm will eventually choose a price that is just below its competitor's price. Given such prices, if the grid is sufficiently fine, then the difference in prices will be less than the cost of search so it cannot be beneficial to search. By proposition 5, consumers gradually adopt higher search thresholds. Price stickiness implies that, over time, prices will almost surely be stuck close together or that the time until a firm relents from the price war will be sufficiently long such that the distribution of consumer thresholds approaches e_L. At this point, the best response correspondence implies that prices should cycle with a maximum price of at least $\hat{p}$. C1 implies that immediately after a firm relents, when the prices are approximately $p^*(e_L)$ and $\xi \in \tilde{P}(e_L)$, the cost of search is sufficiently low so search is beneficial. Again, this will almost surely occur for long enough that the distribution of consumer thresholds approaches e_0, at which point the prices will cycle with a lower bound of approximately $\underline{p}$. Thus, the process cannot converge and these bounds must be approached infinitely many times.

Condition 2 (C2)**.** *If* $p_i = \sup \tilde{P}(e_L)$ *and* $p_{-i} = p_i^*(e_L) = p^*(e_L)$*, then* $c < c^*(p)$*.*

Replacing C1 with the weaker C2 yields the following corollary to Theorem 1.

Corollary 1. *Under A1–A3 and C2, for all* $\varepsilon > 0$*, there exists a grid* G *such that there exists an equilibrium in which*

(i) neither p^t *nor* x^t *converge as* $t \to \infty$*,*

(ii) for all times $t > T$*,* $p_i^t \in [\underline{p} - \varepsilon, \xi^m + \varepsilon]$*, and*

(iii) $(x^t, p_i^t) \to (e_0,\, \xi \leq \underline{p} + \varepsilon)$ *and* $(x^t, p_i^t) \to (e_L,\, \xi' \geq \hat{p} - \varepsilon)$ *infinitely many times.*

Note that C1 and C2 coincide if and only if $\tilde{P}(e_L)$ is a singleton.[29] This result is informative, as Section 3.2.2 shows that an equilibrium exists in which the distribution of consumer thresholds converges under the complement of C1 and all equilibria have this property under the complement of C2. Hence, multiple equilibrium dynamics may exist for a range

[29]The arguments made to prove this corollary are nearly identical to those made in Theorem 1 with one key difference: the grid needs to be chosen such that for some neighborhood $\mathcal{N}(e_L)$ of e_L and all $x \in \mathcal{N}(e_L)$, $\sup \tilde{P}(e_L, G) \in \tilde{P}(x, G)$ and the equilibrium needs to dictate that when a firm chooses a price $p^t \in \tilde{P}(x, G)$, it will choose $p^t = \sup \tilde{P}(e_L, G)$.

of search costs if $\tilde{P}(e_L)$ is not a singleton, though the equilibrium pricing strategy is unique (Proposition 2).

The results of this section somewhat correspond with those of Maskin and Tirole (1988) regarding Edgeworth Cycles, with several important differences. First and foremost, cycles are the only dynamic in the short run and, as we show below, kinked demand equilibria can only emerge in the long run, with cycles defining the short run. Second, the mechanism by which the cycles emerge is different. In Maskin and Tirole, cycles emerge as a best response to a dynamic, forward-looking strategy, whereas cycles in this model emerge due to the presence of boundedly rational consumers which creates positive residual demand for the firm with a higher price. Thus, the cycles in our paper are driven by the demand side of the market, as opposed to Maskin and Tirole, where cycles are driven by the supply side. Third, the characteristics of the cycles differ. In Maskin and Tirole, prices are driven down to marginal cost, and firms relent randomly since relenting yields zero profits in the short run. That is, once price hits marginal cost, firms make zero profits. Because residual demand is zero for the firm with the higher price, relenting leads to zero profits as well. Each firm prefers that its competition relents first, so it can benefit by undercutting. Due to the public nature of relenting, firms employ a mixed strategy once price hits marginal cost. In this setting, relenting is optimal due to non-zero residual demand for the firm with the higher price, allowing identification of not only when the firms will relent, but also which firm will relent.

3.2.2 Kinked Demand Equilibria and the Bertrand and Diamond Paradoxes

When the conditions presented in the previous subsection are violated, the equilibrium converges over time. There are two notions in which the dynamic may converge, depending on the parameters of the model. First, Edgeworth cycles may persist indefinitely, but with the range of the cycles shrinking in the limit. Second, there may be convergence of both prices to the search cost in finite time, though occurrence of this dynamic requires severely restrictive conditions. In either case, the distribution of consumer thresholds converges to e_L. Formally, the conditions considered here are as follows.

Condition 1′ ($C1'$). *If $p_i = \inf \tilde{P}(e_L)$ and $p_{-i} = p_i^*(e_L) = p^*(e_L)$, then $c > c^*(p)$.*

Condition 2′ ($C2'$). *If $p_i = \sup \tilde{P}(e_L)$ and $p_{-i} = p_i^*(e_L) = p^*(e_L)$, then $c > c^*(p)$.*

Condition 3 (C3)**.** *There exists a neighborhood of e_L $\mathcal{N}(e_L)$ such that $\tilde{P}(x) = \{c\}$ for all $x \in \mathcal{N}(e_L)$.*

C2′ is the complement of C2 and C1′ is the complement of C1. While not immediately obvious, C3 is a subcase of C2′.[30] The following theorem characterizes the first type of convergence in which the distribution of consumer thresholds converges to e_L and price cycles indefinitely under a smaller range than Theorem 1.

Theorem 2. *Under A1−A3 and C2′, for all $\varepsilon > 0$, there exists a $\delta > 0$ and $T > 0$ such that if $\|G\| < \delta$, then*

(i) the distribution of consumer thresholds $x^t \to e_L$ as $t \to \infty$,

(ii) for all times $t > T$, $p_i^t \in [p^*(e_L) - \varepsilon, \sup \tilde{P}(e_L) + \varepsilon]$, and

(iii) $p^t \to \xi \leq p^*(e_L) + \varepsilon$ and $p_i^t \to \xi' \geq \hat{p} - \varepsilon$ infinitely many times.

Theorem 2 shows that in the long run, prices cycle indefinitely with a lower bound of $p^*(e_L)$ and an upper bound between $\inf \tilde{P}(e_L)$ and $\sup \tilde{P}(e_L)$. If the firms' prices are close ($|p_1 - p_2| < c$), then consumers do not benefit from searching, so the distribution of consumer search thresholds will tend towards e_L. Given price stickiness and the fact that firms will set prices that are close until one firm relents and raises its price, the firms' prices will almost surely remain close for a sufficiently long period of time such that the distribution of consumer thresholds approaches e_L. At that point, the consumers search sufficiently infrequently such that the gap that emerges following one firm relenting will not be large enough to induce search. As such, the distribution of consumers' thresholds will continue to converge towards e_L and the firms' cycles will remain fixed.

[30] To see why, note that since $p_i^* \leq \inf \tilde{P}(x)$, under C3, $p_i^*(x) = c$ for each firm i and all $x \in \mathcal{N}(e_L)$. It follows that $c^*(p^*(e_L), \sup \tilde{P}(e_L)) = 0$, so $c > 0$ implies C2′.

Corollary 2. *Under A1–A3 and C1′, for all $\varepsilon > 0$, there exists a grid G and a time $T > 0$ such that there exists an equilibrium in which*

(i) the distribution of consumer thresholds $x^t \to e_L$ as $t \to \infty$,

(ii) for all times $t > T$, $p_i^t \in [p^*(e_L) - \varepsilon, \sup \tilde{P}(e_L) + \varepsilon]$, and

(iii) $p^t \to \xi \leq p^*(e_L) + \varepsilon$ and $p_i^t \to \xi' \geq \hat{p} - \varepsilon$ infinitely many times.

Thus C1′ is a sufficient condition for this type of convergence to occur. If $\tilde{P}(e_L)$ is not a singleton, then there is a range of search costs in which both C1′ and C2′ are satisfied. For search costs in that range, there exists both convergent and nonconvergent equilibrium paths. This multiplicity of equilibria can be ruled out by assuming φ is such that the residual profit given $x = e_L$, $\xi D(\xi)(1 - \varphi(\xi - \tau_L))$, is strictly quasiconcave in ξ. The following proposition demonstrates some limiting properties as the grid of consumer thresholds becomes large.

Proposition 6. *Under A1–A3, as $\tau_L \to \xi^m - \inf \operatorname{supp} \varphi$, $p_i^*(e_L) \to \xi^m$ and thus $\sup \tilde{P}(e_L) = p_i^*(e_L)$. Consequently, C2 holds for sufficiently large τ_L. Furthermore for all $\varepsilon > 0$, there exists a $\bar{\tau} > 0$, $\delta > 0$, and $T > 0$ such that if $\tau_L > \bar{\tau}$ and $\|G\| < \delta$, then all equilibria are such that $p_i^t > \xi^m - \varepsilon$ for both firms i and all $t > T$.*

This result follows as a corollary of Theorem 2 and is very similar to the Diamond paradox – equilibrium monopoly pricing with any number of firms – that emerges in the dynamic model of undirected consumer search in Diamond (1971) and the static model of Stahl (1989). The intuition for the result is similar, though the mechanisms by which the equilibrium emerges is different. In this model, the firms will price in such a way that consumers are induced to search less often, and eventually search with sufficiently low probability such that monopoly pricing is optimal. Once reached, search is not beneficial, as both firms set the monopoly price, and so the outcome is stable. In the Diamond-Stahl model of consumer search, the consumers form rational expectations of the firms' pricing distributions, and the only possible equilibrium involves monopoly pricing by the firms and no search by the consumers.[31] With their mechanism, the Diamond paradox emerges because of simultaneous anticipation of this outcome by both firms and consumers.

Proposition 6 posits an equilibrium of similar character to the Diamond paradox (Stahl, 1989) and the kinked demand curve equilibrium of Maskin and Tirole (1988). The difference is in the mechanisms driving the result. In Stahl (1989), convergence to monopoly pricing and no search occurs when the exogenously determined proportion of searchers (with a zero search cost) tends to zero. In Maskin and Tirole (1988), convergence follows from the firm side of the market. As firms become infinitely patient, a Markov perfect equilibrium in which both firms choose the monopoly price can be sustained, as the firms are able to internalize the cost of engaging in a price war. In this paper, the equilibrium follows from demand side characteristics. Moreover, the process by which consumers stop searching is endogenous to the firms' pricing strategies. Thus, a feedback loop occurs where the proximity of the firms' prices during the cycles induce consumers to search less. As consumers search less, the lower bound of the cycles increases (Proposition 4). Furthermore under $C2$, the difference in pricing as the cycles reset is not enough to decrease the consumers thresholds. Thus the process continues until prices approach the monopoly level and the distribution of thresholds approaches e_L.

Lastly, the following theorem demonstrates that under condition $C3$, all equilibria are such that the firms' prices converge in finite time, with the distribution of consumers' converging over time to e_L.

Theorem 3. *Under $A1-A3$ and $C3$, for all $\varepsilon > 0$, there exists a $\delta > 0$ and $T > 0$ such that if $\|G\| < \delta$, then $x^t \to e_L$ and firms will undercut one another until $p_1^t = p_2^t = \xi = c$, and will remain at that price thereafter.*

Under C3, given any distribution of thresholds, the consumers search at any price $\xi > c$ with sufficiently high probability, so the residual profit is always maximized by setting the price equal to the search cost to induce consumers not to search. Thus given some initial prices, the firms engage in a price war until the price is driven to the search cost, and

[31] See Online Appendix B.3 for a full derivation of this result in the context of our model.

the firms never have the incentive to increase their prices. This result can be seen as a Bertrand-like outcome; if we allow the search cost to tend to zero, then the equilibrium will converge to marginal cost pricing, but will not remain as there will be random shifts in the search thresholds, allowing firms to raise their prices and receive positive profits.

Corollary 3. *Under the conditions of Theorem 3, if $c = 0$, then for every $T > 0$, there exists an $\varepsilon > 0$, $\eta > 0$, and $T' > t$ such that $p_i^t \in (0, \eta)$ on $t \in [T', T' + \varepsilon)$.*

That is, the (limiting) Bertrand outcome is Lyapunov stable, though not asymptotically stable.

4 Discussion and Concluding Remarks

By developing a model of undirected consumer search with firm competition, this paper has characterized the market outcomes and developed several empirically testable predictions – specifically the short- and long-run dynamics – resulting from the bounded rationality of agents. The short-run dynamics are characterized by stochastic Edgeworth cycles. For a given distribution of consumer search thresholds, the firm with the larger installed base has the greater incentive to monopolize residual demand rather than continuing the price war. Hence, in situations in which the firms' opportunities to update prices occur more frequently than the consumers' opportunities to reevaluate search decisions, larger firms are more likely to end the reset the cycle by monopolizing residual demand, a prediction consistent with gasoline markets (Noel, 2007b; Atkinson, 2009; Zimmerman et al., 2012; Isakower and Wang, 2014). These cycles persist (aperiodically) in the long run unless consumers can learn to search with sufficient infrequency (almost surely not search). Hence, convergence to the kinked demand equilibria of Maskin and Tirole (1988) and the Diamond paradox occurs only under restricted circumstances.

The model also offers a consistent explanation of the distinct pricing patterns in gasoline markets outlined in Noel (2007a). Noel (2007a) finds three distinct patterns: sticky pricing, Edgeworth price cycles, and cost-based pricing. Each of these regimes is consistent with one of our theorems. The sticky-pricing result is consistent with Theorem 1, Edgeworth cycles with Theorem 2, and cost-based pricing with Theorem 3 (and Proposition 6). Moreover, Corollaries 1 and 2 highlight that multiple equilibrium dynamics can persist in which there is both convergence and non-convergence. Hence, our model proposes a consistent theoretical explanation of these differences. A formal empirical analysis of this mechanism is left for future work.

The bounded rationality of consumers has important implications. The model does not coincide with Maskin and Tirole (1988), Eckert (2003), and Noel (2008) even if search costs are zero or firms are forward looking. Whenever prices equalize, consumers will randomly shift toward higher price thresholds, thereby creating an incentive for firms to raise prices. Both types of equilibria that Maskin and Tirole characterize involve firms setting identical prices at least for some period of time. Thus, the underlying consumer dynamic will induce some change in Maskin and Tirole's results that does not disappear in the limit.

We now conclude with a discussion of the implications of the model and the assumptions surrounding the search costs and decisions. First, a major implication of the results is that the Bertrand paradox does not emerge as an equilibrium outcome as the cost of search tends to zero. This result demonstrates the (asymptotic) instability of the Bertrand paradox as it only occurs under highly specific assumptions and is upset by arbitrary perturbations. Marginal cost pricing cannot emerge with even arbitrarily small search costs in our model because consumers do not benefit from search when prices approach this level. Over time, this will lead to random shifts in the distribution of consumer price thresholds (Proposition 5) that firms can take advantage to obtain positive profits (by A3). Thus, the Bertrand paradox is only a temporary outcome in the singleton case with exactly zero search costs.

Second, the cost of search does not influence any of the qualitative results of the model as long as that cost is nonzero (with a sufficiently fine pricing grid). The cost is only relevant insofar as it determines whether search has a positive or negative expected value, and as long as the price grid is sufficiently fine, it can take on both of these values. The only effect that the cost has on the process is on the rate at which the process evolves. The reason is that the probability that a consumer adopts a strategy is increasing in the payoff from that strategy and decreasing in the payoff of the incumbent strategy. A higher search cost reduces the expected value of search in all states, and thus increases the rate at which consumers are discouraged from search during the price cycles and decreases the rate at which consumers are encouraged to search when the cycle resets.

Though the search cost does not influence the qualitative results, the model offers important insights into the role of the search cost (and its endogenization) in pricing. If higher search costs increase the rate of convergence to higher thresholds and reduce the rate of convergence to lower thresholds, then firms benefit by raising the search cost so that consumers search less for greater periods of time. Furthermore, both the critical judo price and smallest residual maximizer are bounded below by the search cost. Hence, firms can potentially increase their profits via obfuscation (Ellison and Ellison, 2009; Ellison and Wolitsky, 2012) to increase the search cost. When binding (i.e., when the judo price is the search cost or under C3 and the conditions of Theorem 3), firms can strictly benefit from actions that increase the search cost c, provided that the cost of such actions are not too large.

Third, a component that most models of undirected search incorporate is a fraction of "shoppers," consumers that either have zero search cost (and thus optimally search) or are exogenously informed as to the prices set by the firms. As this paper's primary objective is to weaken the assumptions that are often present in these models and examine the impact of bounded rationality on firm and consumer behavior, such an inclusion would be inappropriate for this paper. Nevertheless, it is straightforward to deduce the impact that a fraction of shoppers would have on the model's predictions. These shoppers increase the fraction of consumers that search given a fixed observed price. Thus, the inclusion of shoppers in the model is equivalent to creating a negative bias to the noise in each consumer's price threshold. If the fraction of shoppers is large enough, this will induce the conditions of Theorem 1 and thus perpetual cycling of prices and consumer price thresholds. Otherwise, if this fraction is sufficiently small, then the conditions of Theorem 2 and Proposition 6 may still be satisfied, resulting in a long run convergence to the Diamond paradox.

This result stands in contrast to Stahl (1989), who finds in a static model where consumers form rational expectations about the firms' strategies that the distribution of firm's prices varies continuously between marginal cost pricing (Bertrand paradox) and monopoly pricing

(Diamond paradox) as the fraction of shoppers varies between zero and one.[32] Similarly, this result is in contrast to the theoretical results of Pennerstorfer et al. (2020), who find similar results to Stahl (1989). However, the persistence of price dispersion is consistent with the semiparametric empirical evidence of Pennerstorfer et al. (2020), whereby the degree of price dispersion decreases, but does not disappear as the share of shoppers tends to zero or one.

References

Allison, Blake and Jason Lepore, "A General Model of Bertrand Edgeworth Duopoly," *Working Paper*, 2016.

Anderson, Simon P. and Régis Renault, "Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model," *RAND Journal of Economics*, 1999, *30* (4), 719–735.

Arbatskaya, Maria, "Ordered search," *RAND Journal of Economics*, 2008, *38* (1), 119–126.

Armstrong, Mark, "Ordered consumer search," *Journal of the European Economic Association*, 2017, *15* (5), 989–1024.

Atkinson, Benjamin, "Retail gasoline price cycles: Evidence from Guelph, Ontario using bi-hourly, station-speciific retail price data," *Energy Journal*, 2009, *30*, 85–109.

Baye, Michael R., John Morgan, and Patrick Scholten, "Information, search, and price dispersion," in Terrence Hendershott and Andrew B. Whinston, eds., *Handbooks in Information Systems: Economics and Information Systems*, Vol. 1 2006.

Benabou, Roland and Robert Gertner, "Search with Learning from Prices: Does Increased Inflationary Uncertainty Lead to Higher Markups?," *Review of Economic Studies*, 1993, *60* (1), 69–94.

Bikhchandani, Sushil and Sunil Sharma, "Optimal search with learning," *Journal of Economic Dynamics and Control*, 1996, *20* (1-3), 333–359.

Cabral, Luis and Arthur Fishman, "Business as Usual: A Consumer Search Theory of Sticky Prices and Asymmetric Price Adjustment," *International Journal of Industrial Organization*, 2012, *30* (4), 371–376.

Castanias, Rick and Herb Johnson, "Gas Wars: Retail Gasoline Fluctuations," *Review of Economics and Statistics*, 1993, *75* (1), 171–174.

Cheung, Man-Wah, "Imitative dynamics for games with continuous strategy space," *Games and Economic Behavior*, 2016, *99*, 206–223.

_ **and Jiabin Wu**, "On the probabilistic transmission of continuous cultural traits," *Journal of Economic Theory*, 2018, *174*, 200–323.

Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim, "Consumer search and price competition," *Econometrica*, 2018, *86* (4), 1257–1281.

[32] Equilibrium pricing in Stahl's model is in mixed strategies when the fraction of shoppers is strictly between zero and one.

Dana Jr., James D., "Learning in an Equilibrium Search Model," *International Economic Review*, 1994, *35* (3), 745–771.

Deneckere, Raymond J. and Dan Kovenock, "Price Leadership," *Review of Economic Studies*, 1992, *59* (1), 143–162.

Diamond, Peter, "A Model of Price Adjustment," *Journal of Economic Theory*, 1971, *3* (2), 156–168.

Doyle, Joseph, Erich Muehlegger, and Krislert Samphantharak, "Edgeworth Cycles Revisited," *Energy Economics*, 2010, *32* (3), 651–660.

Eckert, Andrew, "Retail price cycles and the presence of small firms," *International Journal of Industrial Organization*, 2003, *21* (2), 151–170.

_ **and Douglas West**, "Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations," *Journal of Law and Economics*, 2004, *47* (1), 245–273.

Edgeworth, Francis, "A Pure Theory of Monopoly," *Papers Relating to Political Economy*, 1925, *1*, 111–142.

Ellison, Glenn, "Bounded Rationality in Industrial Organization," in Richard Blundell, Whitney K. Newey, and Torsten Persson, eds., *Advances in Economics and Econometrics: Theory and Applications, Ninth World Conference*, Vol. 2, Cambridge University Press, 2006.

_ **and Alexander Wolitsky**, "A search cost model of obfuscation," *RAND Journal of Economics*, 2012, *43* (3), 417–441.

_ **and Sarah Fisher Ellison**, "Search, obfuscation, and price elasticities on the internet," *Econometrica*, 2009, *77* (2), 427–452.

Fershtman, Chaim and Arthur Fishman, "Price cycles and booms: Dynamic search equilibrium," *American Economic Review*, 1992, *82* (5), 1221–1233.

Garcia, Daniel, Jun Honda, and Maarten Janssen, "The double Diamond paradox," *American Economic Journal: Microeconomics*, 2017, *9* (3), 63–99.

Gelman, Judith and Steven Salop, "Judo Economics: Capacity Limitation and Coupon Competition," *The Bell Journal of Economics*, 1983, *14* (2), 315–325.

Isakower, Sean and Zhongmin Wang, "A comparison of regular price cycles in gasoline and liquefied petroleum gas," *Energy Economics*, 2014, *45*, 445–454.

Janssen, Maarten, Alexei Parakhonyak, and Anastasia Parakhonyak, "Non-reservation price equilibria and consumer search," *Journal of Economic Theory*, 2017, *172*, 120–162.

_ **, Paul Pichler, and Simon Weidenholzer**, "Oligopolistic markets with sequential search and production cost uncertainty," *RAND Journal of Economics*, 2011, *42* (3), 444–470.

Lewis, Matthew S., "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market," *Journal of Economics and Management Strategy*, 2011, *20* (2), 409–449.

_ **and Howard P. Marvel**, "When do consumers search?," *Journal of Industrial Economics*, 2011, *59* (3), 457–483.

Maskin, Eric and Jean Tirole, "A Theory of dynamic oligopoly, ii: Kinked demand curves, and Edgeworth cycles," *Working Paper*, 1985.

— and — , "A Theory of Dynamic Oligopoly, II: Kinked Demand Curves, and Edgeworth Cycles," *Econometrica*, 1988, *56* (3), 571–599.

Noel, Michael D., "Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets," *Review of Economics and Statistics*, 2007, *89* (2), 324–334.

— , "Edgeworth price cycles: Evidence from the Toronto retail gasoline market," *Journal of Industrial Economics*, 2007, *55* (1), 69–92.

— , "Edgeworth price cycles and focal points: Computational dynamic Markov equilibria," *Journal of Economics and Management Strategy*, 2008, *17* (2), 345–377.

Osborne, Martin J. and Carolyn Pitchik, "Price competition in a capacity-constrained duopoly," *Journal of Economic Theory*, 1986, *38* (2), 238–260.

Pennerstorfer, Dieter, Phillip Schmidt-Dengler, Nicolas Schutz, Christoph Weiss, and Biliana Yontchevka, "Information and price dispersion: Theory and evidence," *International Economic Review*, 2020, *61* (2), 871–899.

Preuss, Marcel, "Searching, learning, and tracking," *RAND Journal of Economics*, 2023, *54* (1), 54–82.

Rauh, Michael T., "A model of temporary search market equilibrium," *Journal of Economic Theory*, 1997, *77* (1), 128–153.

Rothschild, Michael, "Searching for the lowest price when the distribution of prices is unknown," *Journal of Political Economy*, 1974, *82* (4), 689–711.

Salop, Steven and Joseph Stiglitz, "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, 1977, *44* (3), 493–510.

Sandholm, William, *Population Games and Evolutionary Dynamics*, Cambridge, MA: Massachussetts Institute of Technology, 2010.

Shubik, Martin, "Edgeworth Market Games," in R. Duncan Luce and Albert Tucker, eds., *Contributions to the Theory of Games*, Vol. IV 1959.

Spiegler, Ran, *Bounded Rationality and Industrial Organization*, Oxford Uniersity Press, 2014.

Stahl, Dale O., II, "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 1989, *79* (4), 700–712.

Stigler, George J., "The economics of information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Tappata, Mariano, "Rockets and Feathers: Understanding Asymmetric Pricing," *RAND Journal of Economics*, 2009, *40* (4), 673–687.

Varian, Hal, "A Model of Sales," *American Economic Review*, 1980, *70* (4), 651–659.

Vives, Xavier, "Edgeworth and Modern Oligopoly Theory," *European Economic Review*, 1993, *37* (2-3), 463–476.

Wallner, Klaus, "Sequenial moves and tacit collusion: Reaction-function cycles in a finite pricing duopoly," *Journal of Economic Theory*, 1999, *84* (2), 251–267.

Wang, Zhongmin, "Station Level Gasoline Demand in an Australian Market with Regular Price Cycles," *Australian Journal of Agricultural and Resource Economics*, 2009, *53* (4), 467–483.

Yang, Huanxing and Lixin Ye, "Search with Learning: Understanding Asymmetric Price Adjustments," *RAND Journal of Economics*, 2008, *39* (2), 547–564.

Zimmerman, Paul R., John M. Yun, and Christopher T. Taylor, "Edgeworth price cycles: Evidence from the United States," *Review of Industrial Organization*, 2012, *42*, 297–320.

Appendix

Lemmas 1 and 2 are supporting results used throughout.

Lemma 1. *Under A2, the following statements are true:*

(i) $\bar{\varphi}(\xi, x)$ *is continuous on* $[0, c) \cup (c, \xi^m] \times X$,

(ii) $\bar{\varphi}(\xi, x)$ *is strictly increasing in* ξ *on* $[c, \xi^m]$,

(iii) *if* x *first order stochastically dominates* x', *then* $\bar{\varphi}(\xi, x) < \bar{\varphi}(\xi, x')$ *for all* $\xi \in [c, \xi^m]$.

Proof of Lemma 1.

Proof. Statements (i) and (ii) follow immediately from A2 and the definition of $\bar{\varphi}(\xi, x)$. Item (iii) is trivially true if $\xi \leq c$. We now prove statement (iii). Suppose that $\xi > c$ and x first order stochastically dominates x'. By Abel's lemma, $\bar{\varphi}(\xi, x) = \sum_{k=0}^{L} \varphi(\xi - \tau_k) x_k$ can be rewritten as

$$\varphi(\xi - \tau_L) - \sum_{k=0}^{L-1} \left(\sum_{\ell=0}^{k} x_\ell \right) (\varphi(\xi - \tau_{k+1}) - \varphi(\xi - \tau_k))$$

Hence, for distributions x and x', $\bar{\varphi}(\xi, x) < \bar{\varphi}(\xi, x')$ if and only if

$$\varphi(\xi - \tau_L) - \sum_{k=0}^{L-1} \left(\sum_{\ell=0}^{k} x_\ell \right) (\varphi(\xi - \tau_{k+1}) - \varphi(\xi - \tau_k)) \\ < \varphi(\xi - \tau_L) - \sum_{k=0}^{L-1} \left(\sum_{\ell=0}^{k} x'_l \right) (\varphi(\xi - \tau_{k+1}) - \varphi(\xi - \tau_k)),$$

which simplifies to

$$\sum_{k=0}^{L-1} \left(\sum_{\ell=0}^{k} (x_\ell - x'_\ell) \right) (\varphi(\xi - \tau_{k+1}) - \varphi(\xi - \tau_k)) > 0.$$

By A2, $(\varphi(\xi - \tau_{k+1}) - \varphi(\xi - \tau_k)) < 0$ for all $k = 0, \ldots, L$, so it is sufficient that, for each $k = 0, \ldots, L-1$,

$$\sum_{k=0}^{L-1} \sum_{\ell=0}^{k} (x_\ell - x'_\ell) = \sum_{k=0}^{L} \sum_{\ell=0}^{k} (x_\ell - x'_\ell) < 0,$$

which is true if x exhibits first order stochastic dominance over x'. □

Proof of Proposition 1.

Proof. To demonstrate that $\tilde{P}(x)$ is nonempty, it sufficient to show that for all $x \in X$, $\pi_i^R(\xi, x)$ is upper semicontinuous in ξ. By Lemma 1(i), $\bar{\varphi}(\xi, x)$ is continuous on $[0, c) \cup (c, \xi^m) \times X$. Therefore, it is sufficient verify upper semicontinuity at $\xi = c$. For all $\xi \leq c$, $\bar{\varphi}(\xi, x) = 0$, so $\pi_i^R(\xi, x) = \alpha_i \xi D(\xi)$. For any $\xi > c$, $\pi_i^R(\xi, x) = \alpha_i \xi D(\xi)(1 - \bar{\varphi}(\xi, x)) \leq \alpha_i \xi D(\xi)$. Thus,

$$\limsup_{\xi \to c} \pi_i^R(\xi, x) = \pi_i^R(c, x),$$

so $\pi_i^R(\xi, x)$ is upper semicontinuous and possesses a maximizer on $[0, \xi^m]$. Hence, $\tilde{P}(x)$ is nonempty.

To show that $\sup \tilde{P}(x) \leq \xi^m$, suppose to the contrary that $\xi \in \tilde{P}(x)$ but $\xi > \xi^m$. Then, under A1, $\xi D(\xi) < \xi^m D(\xi^m)$. Judiciously adding zero by adding and subtracting

$$\xi^m D(\xi^m)(1 - \bar{\varphi}(\xi^m, x)) + \xi D(\xi) \bar{\varphi}(\xi^m, x)$$

to this candidate maximal residual profits yields

$$\begin{aligned} \xi D(\xi)(1 - \bar{\varphi}(\xi, x)) = \xi^m D(\xi^m)(1 - \bar{\varphi}(\xi^m, x)) - \xi D(\xi)(\underbrace{\bar{\varphi}(\xi, x) - \bar{\varphi}(\xi^m, x)}_{>0 \text{ by Lemma 1}}) \\ - (\underbrace{\xi^m D(\xi^m) - \xi D(\xi)}_{>0 \text{ by } A1})(1 - \bar{\varphi}(\xi^m, x)) < \xi^m D(\xi^m)(1 - \bar{\varphi}(\xi^m, x)), \end{aligned}$$

a contradiction of ξ as a residual maximizer. Hence, $\sup \tilde{P}(x) \leq \xi^m$. □

Lemma 2. *Under A1 and A2, $c \leq p_i^*(x) < \xi^m$.*

Proof of Lemma 2.

Proof. For $\xi < c$, observe that the front-side profits given prices $p = (\xi, \xi)$ is $\pi^F((\xi, \xi), x) = \alpha_i \xi D(\xi)$. *A*1 implies that the front-side profit is strictly increasing in ξ. Since firm i's residual profits at $p_i = c$ are

$$\pi_i^R(c, x) = \alpha_i c D(c) > \alpha_i \xi D(\xi) = \pi^F((\xi, \xi), x),$$

it must be that $c \leq p_i^*(x)$ as the firm strictly prefers maximizing its residual demand rather than undercutting at any $\xi < c$.

By *A*2, $\bar{\varphi}(\xi^m, x) > 0$. Consider $\xi \in \tilde{P}(x)$ and note that

$$\begin{aligned}\pi_i^F((\xi^m, \xi^m), x) = \xi^m D(\xi^m)(\alpha_i + (1 - \alpha_i)\bar{\varphi}(\xi^m, x)) &> \alpha_i \xi^m D(\xi^m) \\ &\geq \alpha_i \xi D(\xi) \\ &\geq \max_{p_i} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)) \\ &= \max_{p_i}, \pi^R(p_i, x),\end{aligned}$$

where the second inequality follows from $\xi \leq \xi^M$ being a residual maximizer and last inequality follows from Proposition 1. Hence, $p_i^*(x) < \xi^m$. □

The proof of Proposition 2 relies on the following three lemmas.

Lemma 3. *Let* $g_{\omega^*(x)} = p_i^*(x, G)$. *There exists a* $\delta > 0$ *such that if* $\|G\| < \delta$, *then* $g_{\omega^*(x)} < g_M = \xi^m$.

Proof. Suppose to the contrary that $g_{\omega^*(x)} = g_M = \xi^m$ and consider $\xi \in \tilde{P}(x, G)$. The proof proceeds in two cases: $\xi < \xi^m$ and $\xi = \xi^m$.

Case 1: $\xi < \xi^m$. If $\xi < \xi^m$, then $\xi^m = g_M > g_{M-1} \geq \xi$. Under this supposition,

$$g_{M-1} D(g_{M-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_M, x)) < \alpha_i \xi D(\xi)(1 - \bar{\varphi}(\xi, x)), \tag{2}$$

by the definition of the restricted judo price. *A*1 and Lemma 1(ii) imply that $g_{\omega-1} D(g_{\omega-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_\omega, x))$ is weakly increasing in both ω and g_ω, so

$$\begin{aligned}g_{M-1} D(g_{M-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_M, x)) &\geq g_{M-1} D(g_{M-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_{M-1}, x)) \\ &\geq \xi D(\xi)(\alpha_i + (1 - \alpha_i)\bar{\varphi}(\xi, x)) \\ &\geq \alpha_i \xi D(\xi)(1 - \bar{\varphi}(\xi, x)).\end{aligned} \tag{3}$$

By (2) and (3),

$$\alpha_i \xi D(\xi)(1-\bar{\varphi}(\xi,x)) < \alpha_i \xi D(\xi)(1-\bar{\varphi}(\xi,x)),$$

a contradiction.

<u>Case 2</u>: $\xi = \xi^m$. Because $g_{M-1} < g_{\omega^*(x)} = g_M = \xi^m$ by hypothesis, it follows from the definition of the restricted judo price that

$$g_{M-1}D(g_{M-1})(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi^m,x)) < \alpha_i \xi^m D(\xi^m)(1-\bar{\varphi}(\xi^m,x)),$$

which when rearranged yields

$$\xi^m D(\xi^m) - g_{M-1}D(g_{M-1}) > g_{M-1}D(g_{M-1})\frac{\bar{\varphi}(\xi^m,x)}{\alpha(1-\bar{\varphi}(\xi^m,x))}. \tag{4}$$

Suppose that $\|G\| < \delta$ for some $\delta > 0$. Then, $g_{M-1} \geq \xi^m - \delta$, so by $A1$, $g_{M-1}D(g_{M-1}) \geq (\xi^m - \delta)D(\xi^m - \delta)$. Hence, if (4) is satisfied, then

$$\xi^m D(\xi^m) - (\xi^m - \delta)D(\xi^m - \delta) > (\xi^m - \delta)D(\xi^m - \delta)\frac{\bar{\varphi}(\xi^m,x)}{\alpha(1-\bar{\varphi}(\xi^m,x))}. \tag{5}$$

By $A1$, the LHS of (5) is decreasing in δ while the RHS of (5) is increasing in δ. Taking $\delta \to \xi^m$ yields $\xi^m D(\xi^m) > 0$ and taking $\delta \to 0$ yields

$$0 > (\xi^m)D(\xi^m)\frac{\bar{\varphi}(\xi^m,x)}{\alpha(1-\bar{\varphi}(\xi^m,x))}.$$

a contradiction. By continuity, there exists a δ' such that there is a contradiction for all $\delta \leq \delta'$, so for $\|G\| < \delta'$, $g_{\omega^*(x)} = p_i^*(x,G) < g_M$. □

Lemma 4. *There exists a small positive δ such that if $\|G\| < \delta$ and $g_\omega \in R_i(g_\omega, x)$, then $g_\omega \in \tilde{P}(x,G)$.*

Proof. Suppose $g_\omega \in R_i(g_\omega, x)$. Then,

$$\max_{p_i \in G} \alpha_i p_i D(p_i)(1-\bar{\varphi}(p_i,x)) \leq \alpha_i g_\omega D(g_\omega)$$

and

$$g_{\omega-1}D(g_{\omega-1})(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega,x)) \leq \alpha_i g_\omega D(g_\omega), \tag{6}$$

otherwise it must be that $g_\omega \notin R_i(g_\omega, x)$. Define $g_\psi \in G \setminus \{g_M\}$ such that $g_\psi \leq c < g_{\psi+1}$, noting that $\bar{\varphi}(\xi, x) = 0$ for all $\xi \leq g_\psi$. There are two cases to consider: (*i*) $g_\omega \leq g_\psi$ and (*ii*) $g_\omega > g_\psi$.

Case 1: $g_\omega \leq g_\psi$. Then, $\bar{\varphi}(g_\omega, x) = 0$, so

$$\begin{aligned}\max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)) \leq \alpha_i g_\omega D(g_\omega) &= \alpha_i g_\omega D(g_\omega)(1 - \bar{\varphi}(g_\omega, x)) \\ &\leq \alpha_i g_\psi D(g_\psi)(1 - \bar{\varphi}(g_\psi, x)) \\ &\leq \max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)),\end{aligned}$$

confirming the supposition. Thus, if $g_\omega \in R_i(g_\omega, x)$, then $g_\omega \in \tilde{P}(x, G)$.

Case 2: $g_\omega > g_\psi$. Rearranging (6) yields

$$\bar{\varphi}(g_\omega, x) \leq \frac{\alpha_i}{1 - \alpha_i}(g_\omega D(g_\omega) - g_{\omega-1} D(g_{\omega-1})). \tag{7}$$

If $\|G\| < \delta$, then

$$g_\omega D(g_\omega) - g_{\omega-1} D(g_{\omega-1}) \leq g_\omega D(g_\omega) - (g_\omega - \delta) D(g_\omega - \delta)$$

By the uniform continuity of $\xi D(\xi)$ on $[0, \xi^m]$, choose δ such that

$$|\xi D(\xi) - (\xi - \delta) D(\xi - \delta)| < \frac{1 - \alpha_i}{\alpha_i} \varphi(c - \tau_L).$$

Hence, if $\|G\| < \delta$, then (7) implies that $\bar{\varphi}(g_\omega, x) < \varphi(c - \tau_L)$, a contradiction. Thus, it cannot be that $g_\omega > g_\psi$, so the lemma holds. □

Lemma 5. *If* $g_\omega = p_i^*(x, G)$*, then* $g_{\omega-1} \in R_i(g_\omega, x)$*.*

Proof. Suppose that $g_\omega = p_i^*(x, G)$. Because $p_i^*(x, G) \in G$, it must be that

$$g_{\omega-1} D(g_{\omega-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_\omega, x)) = \max_{p_i \in G} \alpha_i D(p_i)(1 - \bar{\varphi}(p_i, x)),$$

or the construction of $p_i^*(x, G)$ would have been perturbed such that $p_i^*(x, G) \notin G$ (see the proof of Proposition 2). By Lemma 4, $g_\omega \in R_i(g_\omega, x)$ implies that $g_\omega \in \tilde{P}(x, G)$, so $g_{\omega-1} \in R_i(g_\omega, x)$. □

Proof of Proposition 2.

Proof. We first prove that, under $A1$ and $A2$, if $\|G\| < \delta$, then $R_i(g_\omega, x)$ is the best response correspondence. Suppose that $\|G\| < \delta < c$. That $p_i^*(x, G) \in G$ follows from $g_0 = 0$, which implies that

$$g_0 D(g_0)(1 - \bar{\varphi}(g_1, x)) = 0,$$

while there is some $g \in G$ with $\in (0, c)$ guaranteeing that

$$\max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)) \geq gD(g) > 0.$$

If

$$g_{\omega^*(x)-1} D(g_{\omega^*(x)-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_{\omega^*(x)}, x)) < \max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)),$$

then perturb the judo price by setting $p_i^*(x, G) = (g_{\omega^*(x)} + g_{\omega^*(x)+1})/2$. This, or any perturbation in $(g_{\omega^*(x)}, g_{\omega^*(x)+1})$ is necessary, as the proposition states that $g_{\omega^*(x)-1}$ is a best response to $p_i^*(x, G)$. In order for this perturbation to be well defined, it must be $\omega^*(x) < M$, which follows from Lemma 3.

$A1$ implies that $\xi D(\xi)(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_\omega, x))$ is strictly increasing in ξ. Thus,

$$R_i(g_\omega, x) \supset \{g_{\omega-1}, g_\omega\} \cup \tilde{P}(x, G).$$

For $R_i(g_\omega, x)$ to be as in the statement of the proposition, two properties must be satisfied: (i) if $g_\omega \in R_i(g_\omega, x)$, then $g_\omega \in \tilde{P}(x, G)$, and (ii) if $g_\omega = p_i^*(x, G)$, then $g_{\omega-1} \in R_i(g_\omega, x)$. The fact that $\tilde{P}(x, G) \subset R_i(g_\omega, x)$ for $g_\omega \leq p_i^*(x, G)$ and $g_{\omega-1} \in R_i(g_\omega, x)$ for $g_\omega > p_i^*(x, G)$ follows directly from the construction of $p_i^*(x, G)$ and the fact that $g_{\omega-1} D(g_{\omega-1})(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_\omega, x))$ is strictly increasing in ω. Lemmas 4 and 5 prove properties (i) and (ii), respectively.

We now prove the two limiting statements. Let G^n be such that $\|G^n\| \to 0$ as $n \to \infty$. We first show that $\max_{g \in G^n} \pi_i^R(g, x) \to \max_\xi \pi_i^R(\xi, x)$, i.e., the constrained maximal residual profits converge to the unconstrained maximal residual profits as the grid becomes arbitrarily fine.

Define $g^n \in G^n$ such that $g^n \to g \in \tilde{P}(x)$ and $g^n < g$. The continuity of π_i^R on $[0, c) \cup (c, \infty) \times X$, coupled with the fact that $\lim_{\xi \to c^-} \pi_i^R(\xi, x) = \pi_i^R(c, x)$ implies that $\pi_i^R(g^n, x) \to$

$\pi_i^R(g, x)$. Thus,

$$\lim_{n\to\infty} \max_{g\in G^n} \pi_i^R(g, x) \geq \max_{\xi} \pi_i^R(\xi, x),$$

while

$$\max_{g\in G^n} \pi_i^R(g, x) \leq \max_{\xi} \pi_i^R(\xi, x)$$

as the unconstrained maximum must be weakly greater than the constrained maximum. Therefore, $\max_{g\in G^n} \pi_i^R(g, x) \to \max_\xi \pi_i^R(\xi, x)$.

Next, define $g^n \in \tilde{P}(x, G^n)$ such that $g^n \to \xi$. By definition, $\pi_i^R(g^n, x) = \max_{g\in G^n} \pi_i^R(g, x)$ and by the preceding argument, $\pi_i^R(g^n, x) \to \max_{p_i} \pi_i^R(p_i, x)$. If $\xi \neq c$, then the continuity of π_i^R guarantees that $\xi \in \tilde{P}(x)$. Now suppose that $\xi = c$. As $\pi_i^R(c, x) = \limsup_{p_i\to c} \pi_i^R(p_i, x)$ and $\max_{g\in G^n} \pi_i^R(g, x) \to \max_{p_i} \pi_i^R(p_i, x)$ implies that $\pi_i^R(c, x) \geq \max_{p_i} \pi_i^R(p_i, x)$, $\xi \in \tilde{P}(x)$, proving statement (i).

We now prove statement (ii). Let $p_i^*(x, G^n) \to \xi^*$. We will show that $\xi^* = p_i^*(x)$. Define $\rho = (\xi^* + p_i^*(x))/2$. The remainder of the proof proceeds in two cases.

<u>Case 1</u>: $\xi^* \geq c$.

Case 1a: $\xi^* < p_i^*(x)$. It follows that there exists a value n^* such that $p_i^*(x, G^n) < \rho$ for all $n > n^*$. Choose such an n^* and note that for all $n > n^*$ and all $g_\omega^n \in G^n \cap (\rho, p_i^*(x))$,

$$\begin{aligned} g_\omega^n D(g_\omega^n)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega^n, x)) &> \rho D(\rho)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\rho, x)) \\ &> \max_{g\in G^n} \pi_i^R(g, x). \end{aligned} \tag{8}$$

The first inequality follows from the fact that $\xi D(\xi)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x))$ is increasing in ξ by $A1$ and Lemma 1 and the second inequality follows from $\rho > p_i^*(x, g)$. Assign $g_\omega^n \in G^n \cap (\rho, p_i^*(x))$ with $g^n \to g < p_i^*(x)$, where this set is nondegenerate for δ sufficiently small. Because $\lim_n \max_{g\in G^n} \pi_i^R(g, x) = \max_\xi \pi_i^R(\xi, x)$, it follows from (8) that

$$gD(g)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g, x)) \geq \max_{\xi} \pi_i^R(\xi, x). \tag{9}$$

However, as $g < p_i^*(x)$,

$$gD(g)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g, x)) < p_i^*(x)D(p_i^*(x))(\alpha_i + (1-\alpha_i)\bar{\varphi}(p_i^*(x), x)).$$

and the definition of the judo price $p_i^*(x)$ implies that,

$$p_i^*(x)D(p_i^*(x))(\alpha_i + (1-\alpha_i)\bar{\varphi}(p_i^*(x), x)) \leq \max_{\xi} \pi_i^R(\xi, x), \tag{10}$$

Collectively, (8)-(10) imply that

$$\begin{aligned}\max_{\xi} \pi_i^R(\xi, x) &\leq gD(g)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g, x)) \\ &< p_i^*(x)D(p_i^*(x))(\alpha_i + (1-\alpha_i)\bar{\varphi}(p_i^*(x), x)) \leq \max_{\xi} \pi_i^R(\xi, x),\end{aligned}$$

a contradiction. Thus, $\xi^* \geq p_i^*(x)$.

Case 1b: $\xi^* > p_i^*(x)$. Let n^* be such that $p_i^*(x, G^n) < \rho$ and note that (by an analogous argument to case 1a) for all $n > n^*$ and all $g_\omega^n \in G^n \cap (p_i^*(x), \rho)$,

$$\begin{aligned}g_\omega^n D(g_\omega^n)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega^n, x)) &< \rho D(\rho)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\rho, x)) \\ &< \max_{g \in G^n} \pi_i^R(g, x).\end{aligned}$$

Now let $g_\omega^n \in G^n \cap (\rho, p_i^*(x))$ with $g^n \to g > p_i^*(x)$. It follows that

$$gD(g)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g, x)) \leq \max_{\xi} \pi_i^R(\xi, x).$$

The definition of $p_i^*(x)$ implies that

$$gD(g)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g, x)) > \max_{g \in G^n} \pi_i^R(g, x),$$

which further implies that $\max_\xi \pi^R(\xi, x) > \max_\xi \pi^R(\xi, x)$, a contradiction.

Case 2: $\xi^* < c$. Recall that $p_i^*(x) \geq c$. Define n^* such that $p_i^*(x, G^n) < \rho$ for all $n > n^*$. Then let $g_\omega^n \in G^n \cap (\rho, p_i^*(x))$ and note that

$$g_{\omega-1}^n D(g_{\omega-1}^n)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega^n, x)) \geq \max_{g \in G^n} \pi_i^R(g, x).$$

Because $g_\omega^n < c$,

$$\begin{aligned}g_{\omega-1}^n D(g_{\omega-1}^n)(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega^n, x)) &= \alpha_i g_{\omega-1}^n D(g_{\omega-1}^n) \\ &< \alpha_i c D(c) \\ &= \pi_i^R(c, x) \\ &\leq \max_{g \in G^n} \pi_i^R(g, x),\end{aligned}$$

a contradiction. Thus, $\xi^* \geq c$ and $\xi^* = p_i^*(x)$, completing the proof. □

Proof of Proposition 3.

Proof. As each $p_i^*(x) \geq c$, the result holds trivially if $p_{-i}^*(x) = c$. Suppose that $p_i^*(x) \in (c, \xi^m)$. By the continuity of $D(\xi)$ and $\bar{\varphi}(\xi, x)$ on $(c, \infty) \times X$, $\pi_{-i}^F(p_{-i}^*(x), x) = \pi_{-i}^R(x)$. That is,

$$p_{-i}^*(x)D(p_{-i}^*(x))(\alpha_{-i} + (1-\alpha_{-i})\bar{\varphi}(p_{-i}^*(x), x)) = \max_{p_{-i}} \alpha_{-i}D(p_{-i})(1 - \bar{\varphi}(p_{-i}, x)).$$

Define the function

$$f(\xi, \alpha) = \xi D(\xi)(\alpha + (1-\alpha)\bar{\varphi}(\xi, x)) - \max_{\zeta} \alpha\zeta D(\zeta)(1 - \bar{\varphi}(\zeta, x)),$$

so that $p_{-i}^*(x)$ is defined by $f(p_{-i}^*(x), \alpha_{-i}) = 0$. Note that $f(\xi, \alpha)$ is strictly increasing in ξ. It therefore suffices to show (by the implicit function theorem) that $f(\xi, \alpha)$ is decreasing in α:

$$\frac{\partial}{\partial \alpha} f(\xi, \alpha) = \xi D(\xi)(1 - \bar{\varphi}(\xi, x)) - \max_{\zeta} \zeta D(\zeta)(1 - \bar{\varphi}(\zeta, x)) \leq 0.$$

If $\alpha_{-i}p_{-i}^*(x)D(p_{-i}^*(x))(1 - \bar{\varphi}(p_{-i}^*(x), x)) = \max_\zeta \alpha\zeta D(\zeta)(1 - \bar{\varphi}(\zeta, x))$, then by definition, $p_{-i}^*(x) \in \tilde{P}(x)$. Thus

$$p_{-i}^*(x)D(p_{-i}^*(x))(\alpha_{-i} + (1-\alpha_{-i})\bar{\varphi}(p_{-i}^*(x), x)) = \alpha_{-i}p_{-i}^*(x)D(p_{-i}^*(x))(1 - \alpha_{-i}\bar{\varphi}(p_{-i}^*(x), x)),$$

which holds only if $\bar{\varphi}(p_{-i}^*(x), x) = 0$, which requires $p_{-i}^*(x) < c$. Because $p_{-i}^*(x) > c$, $\frac{\partial}{\partial \alpha} f(\xi, \alpha) < 0$. Hence, $p_i^*(x) > p_{-i}^*(x)$. □

Proof of Proposition 4.

Proof. Suppose that x first order stochastically dominates x'. Lemma 1(iii) implies that $\bar{\varphi}(\xi, x) \leq \bar{\varphi}(\xi, x')$ for all ξ. Thus,

$$\xi D(\xi)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x)) \leq \xi D(\xi)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x'))$$

and

$$\max_{p_i} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)) \geq \max_{p_i} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(\xi, x')).$$

If

$$\xi D(\xi)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x')) \leq \max_{p_i} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x')), \tag{11}$$

then

$$\xi D(\xi)(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x)) \leq \max_{p_i} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)). \tag{12}$$

Therefore

$$\sup\left\{\xi \leq \xi^m : \xi D(\xi)\left(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x')\right) < \max_{p_i} \alpha_i p_i D(p_i)\left(1 - \bar{\varphi}(p_i, x')\right)\right\}$$
$$\leq \sup\left\{\xi \leq \xi^m : \xi D(\xi)\left(\alpha_i + (1-\alpha_i)\bar{\varphi}(\xi, x)\right) < \max_{p_i} \alpha_i p_i D(p_i)\left(1 - \bar{\varphi}(p_i, x)\right)\right\},$$

and so by definition, $p_i^*(x) \geq p_i^*(x')$.

Now consider the case in which prices are constrained to a grid G with $\|G\| < \delta$, where $\delta > 0$ is sufficiently small such that each firm's best response correspondence is as in Proposition 2. Since $\bar{\varphi}(\xi, x) \leq \bar{\varphi}(\xi, x')$, it follows that for all $g_\omega \in G$,

$$g_{\omega-1} D(g_{\omega-1})(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega, x)) \leq g_{\omega-1} D(g_{\omega-1})(\alpha_i + (1-\alpha_i)\bar{\varphi}(g_\omega, x'))$$

and

$$\max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(p_i, x)) \geq \max_{p_i \in G} \alpha_i p_i D(p_i)(1 - \bar{\varphi}(\xi, x')).$$

Thus, by an identical argument, the unperturbed critical price $p_i^*(x, G) \geq p_i^*(x', G)$. Furthermore, the same argument is valid given strict inequalities in (11) and (12). Therefore, if $p_i^*(x', G)$ is perturbed as in the proof of Proposition 2, then either $p_i^*(x, G) > p_i^*(x', G)$ or $p_i^*(x, G)$ is also perturbed and the inequality holds. □

Proof of Proposition 5.

Proof. First, note that the case in which $c = c^*(p)$ can be ignored, as the grid may be perturbed so that no prices satisfy this relationship. Let $p^t = p$ for all $t \in [T, T + \varepsilon)$. By the definition of $c^*(p)$, search has a negative expected payoff if and only if $c > c^*(p)$. Consequently, a consumer's expected payoff is strictly increasing in her threshold τ when $c > c^*(p)$ as the probability of search is decreasing in the threshold τ.

Under the best response dynamic, if $c > c^*(p)$, then x_L^t is increasing and x_k^t is decreasing for all $k \neq L$. By an analogous argument, if $c < c^*(p)$, then x_0^t is increasing and x_k^t is decreasing for all $k \neq 0$. Thus, the result holds under the best response dynamic.

Now, consider the imitation dynamic. Suppose $c > c^*(p)$. Let $r_{k\ell} = r_{k\ell}(E[\upsilon|p^t, \tau_k], E[\upsilon|p^t, \tau_\ell])$. Then by $A3$, $r_{k\ell} > r_{\ell k}$ if and only if $k < \ell$ because search is not profitable when $c > c^*(p)$. Recall that the net flow of x^t under the imitation dynamic is

$$\begin{aligned} \dot{x}_k^t &= \sum\nolimits_\ell x_\ell^t \rho_{\ell k} - x_k^t \sum\nolimits_\ell \rho_{k\ell} \\ &= x_k^t \sum\nolimits_\ell x_\ell^t (r_{\ell k} - r_{k\ell}). \end{aligned} \tag{13}$$

We now show that for all $t, t' \in [T, T+\varepsilon)$ with $t > t'$,

$$\sum\nolimits_{\ell=0}^k x_\ell^t \leq \sum\nolimits_{\ell=0}^k x_\ell^{t'}$$

for all $k = 0 : L$. It is sufficient to show that, for all $a = 0 : L$,

$$\sum\nolimits_{k=0}^a \dot{x}_k^t \leq 0.$$

For all $a = 0 : L$ and by (13),

$$\begin{aligned} \sum\nolimits_{k=0}^a \dot{x}_k^t &= \sum\nolimits_{k=0}^a x_k^t \sum\nolimits_\ell x_\ell^t (r_{\ell k} - r_{k\ell}) \\ &= \underbrace{\sum\nolimits_{k=0}^a \sum\nolimits_{\ell=0}^a x_k^t x_\ell^t (r_{\ell k} - r_{k\ell})}_{=0} + \sum\nolimits_{k=0}^a \sum\nolimits_{\ell=a+1}^L x_k^t x_\ell^t (r_{\ell k} - r_{k\ell}) \\ &= \sum\nolimits_{k=0}^a \sum\nolimits_{\ell=a+1}^L x_k^t x_\ell^t (r_{\ell k} - r_{k\ell}). \end{aligned}$$

As argued above, $r_{\ell k} > r_{k\ell}$ for all $\ell = a+1 : L$ and all $k = 0 : a < \ell$. Hence,

$$\sum\nolimits_{k=0}^a \dot{x}_k^t = \sum\nolimits_{k=0}^a \sum\nolimits_{\ell=a+1}^L x_k^t x_\ell^t (r_{\ell k} - r_{k\ell}) < 0,$$

completing this case. The proof for $c < c^*(p)$ proceeds identically. □

The following Lemma is used to prove Theorem 1.

Lemma 6. *Under A1, A2, and C1, there exists a* $\delta > 0$ *such that if* $\|G\| < \delta$*, then there exists a neighborhood* $\mathcal{N}(e_L)$ *of* e_L *such that* $c < c^*(p)$ *for all* $x \in \mathcal{N}(e_L)$*, where* $p_i = \inf \tilde{P}(x, G)$ *and* $p_{-i} = p_i^*(x, G) = p^*(x, G)$ *is the critical grid-constrained judo price.*

Proof. First, we prove that there exists a neighborhood $\mathcal{N}(e_L)$ of e_L such that $\tilde{P}(x,G) = \tilde{P}(e_L,G)$ for all $x \in \mathcal{N}(e_L)$. Take $g \in \tilde{P}(e_L,G)$ and define

$$\Delta = \pi_i^R(g,e_L) - \max_{g' \in G \setminus \tilde{P}(e_L,G)} \pi_i^R(g',e_L)$$

as the difference between the maximal constrained residual profits and the second-best. Because π_i^R is continuous in x, there exists a neighborhood $\mathcal{N}(e_L)$ of e_L such that

$$\pi_i^R(g,x) > \pi_i^R(g,e_L) - \frac{\Delta}{2}$$
$$\pi_i^R(g',x) < \pi_i^R(g',e_L) + \frac{\Delta}{2}$$

for all $g' \in G \setminus \tilde{P}(e_L,G)$ and all $x \in \mathcal{N}(e_L)$. Therefore, $\pi_i^R(g,x) > \pi_i^R(g',x)$ for all $x \in \mathcal{N}(e_L)$ and $g' \in G \setminus \tilde{P}(e_L,G)$. Hence, $g \in \tilde{P}(x,G)$ and $\tilde{P}(x,G) = \tilde{P}(e_L,G)$. As e_L first order stochastically dominates all $x \in X$, $p^*(x,G) \leq p^*(e_L,G)$ for all $x \in X$ by Proposition 4.

We now prove that, for all $x \in \mathcal{N}(e_L)$, $c < c^*(p)$ when $p_i = \inf \tilde{P}(x,G)$ and $p_{-i} = p_i^*(x,G) = p^*(x,G)$. By $C1$, $c < \min\{c^*((p^*(e_L), \inf \tilde{P}(e_L))), c^*((\inf \tilde{P}(e_L), p^*(e_L)))\}$. Recall that $c^*(p)$ is defined by

$$u(D(\min p), \min p) - c^*(p) = \alpha u(D(p_1), p_1) + (1-\alpha) u(D(p_2), p_2). \tag{14}$$

The continuity of u and D imply that (14) is decreasing in $\min p$ and increasing in $\max p$. To illustrate this relationship, suppose that $p_1 \leq p_2$. Then, (14) can be rearranged as

$$c^*(p) = (1-\alpha)(u(D(p_1), p_1) - u(D(p_2), p_2)),$$

which by inspection is decreasing in $p_1 = \min p$ and increasing in $p_2 = \max p$.

Set $\varepsilon' = c^*(p) - c$ for the prices given in the statement of the Lemma and define $\underline{x}$ in the closure of $\mathcal{N}(e_L)$ such that $\underline{x}$ is first order stochastically dominated by all $x \in \mathcal{N}(e_L)$. Fix $\varepsilon > 0$. Define $\delta(\varepsilon)$ such that if both $|p_i - p_i'| \leq \delta(\varepsilon)$ and $|p_{-i} - p_{-i}| \leq \delta(\varepsilon)$, then $|c^*(p) - c^*(p')| < \varepsilon$. By Proposition 2, δ' can be chosen such that if $\|G\| < \delta'$, then $|p^*(\underline{x}, G) - p^*(e_L)| \leq \delta(\varepsilon)$ and $\min \tilde{P}(e_L, G) \geq \inf \tilde{P}(e_L) - \delta(\varepsilon)$. By Proposition 4, if $|p^*(\underline{x}, G) - p^*(e_L)| \leq \delta(\varepsilon)$, then $|p^*(e_L, G) - p^*(e_L)| < \delta(\varepsilon)$. Because $c^*(p)$ is decreasing in $\min p$ and $\min \tilde{P}(\underline{x}, G) = \min \tilde{P}(e_L, G) \geq \inf \tilde{P}(e_L) - \delta(\varepsilon)$, a well-defined neighborhood $\mathcal{N}(e_L)$ exists for every ε' such that

$$\min\{c^*(p^*(\underline{x}, G), \min \tilde{P}(\underline{x}, G)), c^*(\min \tilde{P}(\underline{x}, G), p^*(\underline{x}, G))\} \\ > \min\{c^*((p^*(e_L), \inf \tilde{P}(e_L))), c^*((\inf \tilde{P}(e_L), p^*(e_L)))\} - \varepsilon. \quad (15)$$

Assigning $\varepsilon < \min\{c^*((p^*(e_L), \inf \tilde{P}(e_L))), c^*((\inf \tilde{P}(e_L), p^*(e_L)))\} - c$, which is well defined by $C1$, implies that $c < \min\{c^*(p^*(x, G), \min \tilde{P}(x, G)), c^*(\min \tilde{P}(x, G), p^*(x, G))\}$ for all $x \in \mathcal{N}(e_L)$ and G such that $\|G\| < \delta'$. □

Proof of Theorem 1.

Proof. Suppose that $\|G\| < \delta < c$, where δ is such that best response correspondence for each firm is as stated in Proposition 2.

First, we prove statement (i) (the prices do not converge). To the contrary, suppose that $p^t \to (g_\omega, g_{\omega'})$. Without loss of generality, assume that $g_\omega \leq g_{\omega'}$. By Proposition 2, for some time $T > 0$, it must that (i) $g_\omega \leq p_2^*(x^t, G)$ for all $t > T$, (ii) $g_{\omega'} \in \tilde{P}(x, G)$, and (iii) $g_\omega \in \{g_{\omega'}, g_{\omega'-1}\}$. As $|p_1^t - p_2^t| \leq \delta < c$ for all $t > T$, Proposition 5 implies that $x^t \to e_L$ as $t \to \infty$ because at such prices, $\dot{x}_L^t \to 0$ if and only if $x_L \to 1$. Consider two cases.

<u>Case 1</u>: $g_\omega = g_{\omega'}$. In this case, it must be that $g_\omega \leq p_1^*(x^t, G)$ for all $t > T$. Because $p_i^*(x, G) \leq \min \tilde{P}(x, G)$ for all $x \in X$, $p^*(x^t, G) = \min \tilde{P}(x^t, G)$ for all $t > T$. Hence,

$$\min\{c^*((p^*(x^t, G), \min \tilde{P}(x^t, G))), c^*((\min \tilde{P}(x^t, G), p^*(x^t, G)))\} = 0 < c$$

for all $t > T$, which contradicts Lemma 6.

<u>Case 2</u>: $g_\omega = g_{\omega'-1}$. In this case, it must be that $g_{\omega'} \geq p_1^*(x^t, G)$ for all $t > T$. It follows that either $p^*(x^t) = g_\omega$ or that $p^*(x^t) = g_{\omega'}$ for all $t > T$. Either way,

$$\min\{c^*((p^*(x^t, G), \min \tilde{P}(x^t, G))), c^*((\min \tilde{P}(x^t, G), p^*(x^t, G)))\} \leq \delta < c$$

for all $t > T$, which contradicts Lemma 6. Therefore, the prices do not converge.

Next, we prove statement (ii) (in the limit, prices are bounded). By Proposition 2, given a distribution x^t, neither firm i will ever choose a price $p_i^t < g_{\omega-1}$ when $g_\omega = p_i^*(x, G)$.

Thus, by Propositions 2 and 4, neither firm i will ever choose a price below $p_i^t < g_{\omega-1}$ when $g_\omega = p_i^*(e_0, G)$. Set $g_\omega = p_i^*(e_0, G)$ and suppose that $p_i^*(e_0, G) = p^*(e_0, G)$. If $p_i^t = g_{\omega'}$, then $p_{-i}^t \geq g_{\omega'-1}$. As firm i will never choose a price $p_i^t < g_{\omega-1}$, then there will be some time T such that firm i eventually chooses a price $p_i^t \geq g_{\omega-1}$. It follows that for all times $t > T$, $p_i^t \geq g_{\omega-1}$ and $p_{-i}^t \geq g_{\omega-2}$. Therefore, Proposition 2 guarantees that for sufficiently small $\delta > 0$, if $\|G\| < \delta$, then $|p_i^*(e_0, G) - p_i^*(e_0)| < \varepsilon$, and so given such a δ, $p_i^t \geq \underline{p} - \varepsilon$ for both firms i and all $t > T$. That $p_i^t \leq \xi^m$ for all $t > T$ follows directly from $A1$ (and $\max G = \xi^m$).

Lastly, we prove statement (iii) (infinite cycles between e_0 and e_L). To prove this statement, it is sufficient to show that for any $\varepsilon' > 0$ and any neighborhoods $\mathcal{N}(e_0)$ of e_0 and $\mathcal{N}(e_L)$ of e_L, for all T, there is a positive probability that (a) $x^t \in \mathcal{N}(e_L)$ for some $t > T$, (b) $p_i^{t'_n} \to \xi' \geq \hat{p} - \varepsilon$ for some $t > T$, (c) $x^t \in \mathcal{N}(e_0)$ for some $t > T$, and (d) $p_i^{t_n} \to \xi \leq \underline{p} + \varepsilon$ for some $t > T$. We will jointly demonstrate (a) and (b) followed by (c) and (d).

Proposition 2 dictates that at some time t, firms will set their prices such that $p_i^t = g_\omega$ and $p_{-i}^t = g_{\omega-1}$ for some ω. Given such prices, Proposition 5 implies that the distribution of consumer thresholds will be shifting toward e_L. Under the best response dynamic

$$\dot{x}_L^t = 1 - x_i^t,$$

while under the imitation dynamic

$$\dot{x}_L^t = x_L^t \sum_{\ell=0}^{L} x_\ell^t (r_{\ell L} - r_{L\ell}).$$

Given either dynamic, if the prices are fixed at $p_i^t = g_\omega$ and $p_{-i}^t = g_{\omega-1}$ for some ω, then $\dot{x}_L^t \to 0$ if and only if $x_L \to 1$. If these prices were to remain fixed, then $x^t \to e_L$.

Let $\mathcal{N}(e_L)$ be a neighborhood of e_L. Define

$$\tilde{T} = \sup_{p \in G^2} \sup_{x^T \in X} \inf\{t \geq 0 : x^{T+t} \in \mathcal{N}(e_L) \text{ given } p\}.$$

Given any time T, any prices g_ω and $g_{\omega-1}$, and any $x^T \in X$, it follows that if $p_i^t = g_\omega$ and $p_{-i}^t = g_{\omega-1}$ for all $t > T$, then $x^{T+t} \in \mathcal{N}(e_L)$ for all $t > \tilde{T}$. Given the stickiness of pricing, for any prices $p_i^T = g_\omega$, $p_{-i}^T = g_{\omega-1}$, and any $t' > \tilde{T}$, there is a positive probability that $p_i^{T+t} = g_\omega$ and $p_{-i}^{T+t} = g_{\omega-1}$ for all $t < t'$. Thus, there is a positive probability that $x^t \in \mathcal{N}(e_L)$ for some $t > T$.

Choose $\mathcal{N}(e_L)$ such that for all $x \in \mathcal{N}(e_L)$, $p_i^*(x, G) = p_i^*(e_L, G)$ for each firm i and $\tilde{P}(x, G) \subseteq \tilde{P}(e_L, G)$ (By Lemma 6, this neighborhood is well defined). If $x^T \in \mathcal{N}(e_L)$, define

$$\tilde{T}(x^T) = \inf_{p \in G^2} \inf\{t \geq 0 : x^{T+t} \notin \mathcal{N}(e_L) \text{ given } p\}.$$

Then, given any time T such that $x^T \in \mathcal{N}(e_L)$, let Q denote the maximum number of sequential price changes by the firms according to the best response correspondence that are necessary to reach a pair of prices such that $p_i \in \tilde{P}(x, G)$ and $p_{-i} = g_{\omega-1}$, where $g_\omega = p_{-i}$. There is a positive probability that the firms are able to make Q sequential price changes in the time interval $(T, T + \tilde{T}(x^T))$. If δ is chosen sufficiently small so that $\inf \tilde{P}(e_L, G) \geq \inf \tilde{P}(e_L) - \varepsilon$, then there is a positive probability that for either firm i, for some time $t > T$, $p_i^t \geq \hat{p} - \varepsilon$.

Next, using Lemma 6, choose δ and $\mathcal{N}(e_L)$ such that if $\|G\| < \delta$ and $x \in \mathcal{N}(e_L)$, then $c < c^*(p)$, where $p_i = \inf \tilde{P}(x, G)$ and $p_{-i} = p_i^*(x, G) = p^*(x, G)$. Given $x^T \in \mathcal{N}(e_L)$, there will be some $t > T$ where p^t is such that $c < c^*(p^t)$. Let $\mathcal{N}(e_0)$ be a neighborhood of e_0 and define

$$\tilde{T}_0 = \sup_{p \in G^2} \sup_{x^T \in X} \inf\{t \geq 0 : x^{T+t} \in \mathcal{N}(e_0) \text{ given } p\}.$$

Then given $c < c^*(p^T)$, for any $t' > \tilde{T}_0$ there is a positive probability that the firms prices remain fixed for all $t \in [T, t)$, and thus that $x^{T+t'} \in \mathcal{N}(e_0)$.

Finally, choose $\mathcal{N}(e_0)$ and δ such that for all $x \in \mathcal{N}(e_0)$, $p^*(x, G) = p^*(e_0, G) < p^*(e_0) + \varepsilon$. Then define for all $x^T \in \mathcal{N}(e_0)$

$$\tilde{T}_0(x^T) = \inf_{p \in G^2} \inf\{t \geq 0 : x^{T+t} \notin \mathcal{N}(e_L) \text{ given } p\}$$

and let Q_0 be the maximum number of sequential price changes by the firms according to the best response correspondence that are necessary to reach a pair of prices such that $p_i = g_\omega = p^*(x, G)$ and $p_{-i} \in \{g_{\omega-1}, g_{\omega+1}\}$. Given any $x^T \in \mathcal{N}(e_0)$ there is a positive probability that the firms are able to make Q_0 sequential price changes in the time interval $(T, T+\tilde{T})_0(x^T))$, and thus that $p_i^t < p^*(e_L, G) + \varepsilon$ for each firm i for some $t > T$. $\square$

The following Lemma is used to prove Theorem 2.

Lemma 7. *Under A1, A2, and C2', there exists a $\delta > 0$ such that if $\|G\| < \delta$, then there exists a neighborhood $\mathcal{N}(e_L)$ of e_L such that $c > c^*(p)$ for all $x \in \mathcal{N}(e_L)$, where $p_i = \sup \tilde{P}(x, G)$ and $p_{-i} = p_i^*(x, G) = p^*(x, G)$.*

Proof. The proof of this lemma mirrors that of Lemma 6. By an analogous argument to Lemma 6 there exists a neighborhood $\mathcal{N}(e_L)$ of e_L such that $\tilde{P}(x, G) = \tilde{P}(e_L, G)$. As in the proof of Lemma 6, e_L first order stochastically dominates all $x \in X$, so $p^*(x, G) \leq p^*(e_L, G)$ for all $x \in X$.

We now prove that, for all $x \in \mathcal{N}(e_L)$, $c > c^*(p)$ when $p_i = \max \tilde{P}(x, G)$ and $p_{-i} = p_i^*(x, G) = p^*(x, G)$. By $C2$, $c > \min\{c^*((p^*(e_L), \sup \tilde{P}(e_L))), c^*((\sup \tilde{P}(e_L), p^*(e_L)))\}$. Fix $\varepsilon > 0$. Again, define $\delta(\varepsilon)$ such that if both $|p_i - p_i'| \leq \delta(\varepsilon)$ and $|p_{-i} - p_{-i}| \leq \delta(\varepsilon)$, then $|c^*(p) - c^*(p')| < \varepsilon$. By Proposition 2, δ' can be chosen such that if $\|G\| < \delta'$, then $\max \tilde{P}(e_L, G) \geq \sup \tilde{P}(e_L) - \delta(\varepsilon)$. Because $c^*(p)$ is decreasing in $\min p$ and thus $p^*(x, G)$ while $\max \tilde{P}(x, G) = \max \tilde{P}(e_L, G)$ for all $x \in \mathcal{N}(e_L)$,

$$\min\{c^*(p^*(x, G), \max \tilde{P}(x, G)), c^*(\max \tilde{P}(x, G), p^*(x, G))\}$$
$$< \{c^*((p^*(e_L), \sup \tilde{P}(e_L))), c^*((\sup \tilde{P}(e_L), p^*(e_L)))\} + \varepsilon.$$

Assigning $\varepsilon < c - \min\{c^*((p^*(e_L), \sup \tilde{P}(e_L))), c^*((\sup \tilde{P}(e_L), p^*(e_L)))\}$, which is well defined by $C2'$, implies that $c > \min\{c^*(p^*(x, G), \max \tilde{P}(x, G)), c^*(\max \tilde{P}(x, G), p^*(x, G))\}$ for all $x \in \mathcal{N}(e_L)$ and G such that $\|G\| < \delta'$. $\square$

Proof of Theorem 2.

Proof. Suppose that $\|G\| < \delta < c$, where δ is such that the best response correspondences are as stated in Proposition 2. We first show that $x^t \to e_L$. By Proposition 2, there exists a time $T \geq 0$ such that $p_i = g_\omega$ and $p_{-i} \in \{g_\omega, g_{\omega \pm 1}\}$. Hence, $|p_i - p_{-i}| < c$. Given price stickiness, there is positive probability that these prices will remain in this interval until some time $T' > T$. By proposition 5, for a sequence of times $\{t_n\} \subset [T', T)$, x^{t_n} first order stochastically dominates $x^{t_{n'}}$ for all $n' > n$. As e_L first order stochastically dominates all $x \in X$, there exists a $t' \in [T', T)$ such that a state $x^{t'} \in \mathcal{N}(e_L)$ is reached with positive probability, where $\mathcal{N}(e_L)$ is chosen such that Lemma 7 applies. Lemma 7 and Proposition 5 imply that e_L is an absorbing state. Hence, $x^t \to e_L$.

Statement (ii) follows from an identical argument to the proof of statement (ii) of Theorem 1 but fixing $x^t \in \mathcal{N}(e_L)$. Statement (iii) then follows from the best response correspondences of Proposition 2. □

Proof of Proposition 6.

Proof. As $\tau_L \to \xi^m - \inf \operatorname{supp} \varphi$, $\xi^m \leq \tau_L + \sigma^t$ for all σ^t. Hence, $\varphi(\xi^m - \tau_L) \to 0$ as $\tau_L \to \xi^m - \inf \operatorname{supp} \varphi$. Given state $x = e_L$ and $\tau_L \to \xi^m - \inf \operatorname{supp} \varphi$,

$$\pi^R(\xi, e_L) = \alpha_i \xi D(\xi)\left(1 - \varphi(\xi - \tau_L)\right) \leq \alpha_i \xi^m D(\xi^m) = \pi_i^R(\xi^m, e_L)$$

Hence, $\tilde{P}(e_L) = \{\xi^m\}$. Therefore, the judo price is

$$p_i^*(e_L) = \sup\left\{\xi \leq \xi^m : \xi D(\xi)\left(\alpha_i + (1 - \alpha_i)\varphi(\xi - \tau_L)\right) < \alpha_i \xi^m D(\xi^m)\right\}.$$

Evaluating firm i's front-side profits as $p_i \to \xi^m$ and $p_{-i} = \xi^m$ yields

$$\alpha_i \xi^m D(\xi^m) = \xi^m D(\xi^m)(\alpha_i + (1 - \alpha_i)\varphi(\xi^m - \tau_L)).$$

Hence, $p_i^*(e_L) = \xi^m = \sup \tilde{P}(e_L)$. That $C2$ holds for sufficiently large τ_L follows immediately.

Suppose that $\|G\| < \delta < c$ so that each firm's best response correspondence is as in Proposition 2. By continuity, for every $\varepsilon' > 0$, There exists a neighborhood $\mathcal{N}(\xi^m - \inf \operatorname{supp} \varphi)$ such that if $\tau_L \in \mathcal{T}(\xi^m - \inf \operatorname{supp} \varphi)$, then $p_i^*(e_L) > \xi^m - \varepsilon$. Let $\bar{\tau} = \inf \mathcal{N}(\xi^m - \inf \operatorname{supp} \varphi)$. Then, for all $\tau > \bar{\tau}$, $p_i^*(e_L) > \xi^m - \varepsilon$. For a sufficiently small δ, $|p_i^*(e_L, G) - p_i^*(e_L)|$ is sufficiently small such that $p^*(e_L, G) \geq p_i^*(e_L, G) > \xi^m - \varepsilon$.

As $A1$-$A3$ and $C2'$ are all satisfied, Theorem 2 implies that $x^t \to e_L$. Hence, there exists a time $T \geq 0$ such that for all $t > T$, $p_i^t > \xi^m - \varepsilon$. □

Proof of Theorem 3.

Proof. Suppose $A1$-$A3$ and $C3$ are satisfied and set δ sufficiently small such that the best response correspondences are as in Proposition 2. By $C3$, given a neighborhood $\mathcal{N}(e_L)$, $p^*(x) = \sup \tilde{P}(x) = c$ for all $x \in \mathcal{N}(e_L)$. As $C3$ implies $C2$, there exists a time $T \geq 0$ such that $x^t \in \mathcal{N}(e_L)$ for all $t \geq T$. For any prices $p^t = (p_i^t, p_{-i}^t) \geq (c, c)$, Proposition 2 implies that the two firms will undercut each other at each revision opportunity until $p^t = (c, c)$. Because $\tilde{P}(x) = \{c\}$ for all $x \in \mathcal{N}(e_L)$, there exists a time $T' \geq T$ such that $p^t = (c, c)$ for all $t > T'$. □