

Dynamic Price Competition and Evolutionary Behavior with Search*

Blake Allison Michael Sacks
Emory University Clarkson University

Abstract

This paper develops a model of dynamic pricing with boundedly rational consumer search. Consumers observe one price, can choose to 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 dynamically tuned reservation price. Firms optimize given the current distribution of reservation prices. Short run pricing is characterized by Edgeworth cycles. Rather than resetting as a mixed strategy, each cycle resets when the firm with the larger installed base relents by monopolizing its residual demand. The both the lower and upper bounds of the cycles move with consumer search, though not necessarily in the same direction. Whether these cycles persist in the long run or the pricing dynamic converges to Diamond, Bertrand, or a kinked-demand equilibrium depends on exogenous conditions that we outline.

October 29, 2025

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: David D. Reh 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 additionally grateful for comments by seminar participants UC Irvine.

1 Introduction

The dynamic pricing and consumer search literatures introduced several well-known equilibria and paradoxes into industrial organization. From the consumer search literature (Diamond, 1971; Stahl, 1989), the Diamond paradox emerged.¹ From the dynamic pricing literature, three patterns of interest emerged: the Bertrand paradox, the kinked-demand equilibrium, and Edgeworth price cycles (Maskin and Tirole, 1988). By integrating boundedly rational consumer search dynamics with price competition, this paper shows that Edgeworth cycles offer a link between the established equilibrium patterns. Short-run pricing is characterized by Edgeworth cycles, which persist in the long run if consumers are not sufficiently discouraged from search. Otherwise, there is convergence via contracting Edgeworth cycles to a kinked demand equilibrium, either at the monopoly price (the Diamond paradox) or the search cost.

Our model imports tools from evolutionary game theory and 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 into a continuous time setting in which both consumers and firms are boundedly rational.² With a few notable exceptions, 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 (a reservation price equilibrium). The formation of such expectations is generally untenable as consumers often lack the necessary knowledge of the firms’ production processes or pricing policies. A literature has also emerged studying non-reservation price equilibria.³

¹The 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).

²See, e.g., Conlisk (1996), Blinder et al. (1998), Ellison (2006), Spiegel (2014), and Heidhues and Kőszegi (2018) on the use of bounded rationality in industrial organization.

³Instead 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 con-

Instead, we model consumers as making their decisions according to a simple rule of thumb and periodically adjusting their guiding rule via processes that do not require informational or cognitive burdens. 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 individual-specific random shock. Consumers then update their thresholds by imitating the better performing strategies of their peers. This contrasts our notion of reservation price with the typical notion (see, e.g., Anderson and Renault, 2018), where reservation prices are formed via rational expectations over the distribution of prices. Over time, consumers adjust their thresholds according by imitating better performing strategies of their peers.

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. With this limited knowledge, the firms choose their prices to maximize short term profits, though the main results are (qualitatively) not sensitive to this assumption.⁴ 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. Each consumer is endowed with unit demand and will either purchase one unit from the observed firm or engage in search and purchase one unit from the firm with the lowest price.⁵

In this framework, the consumer dynamic can be simply characterized: consumers search

sumers form expectations of prices based on the observed prices during previous purchases. Parakhonyak and Sobolev (2015) shows that consumers employ a random stopping rule and price dispersion persists when do not have prior beliefs on the distribution of prices. 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.

⁴This assumption allows us to isolate the short-run market dynamics that arise when firms react to current profit opportunities. A fully forward-looking formulation would require dynamic optimization over the evolutionary path of consumer search thresholds, creating a high-dimensional state space that is analytically intractable (Weintraub et al., 2008). Modeling firms as myopic therefore clarifies the demand-side mechanism driving non-stationary Edgeworth cycles. This approach follows a growing literature on bounded rationality in industrial organization that highlights limited foresight and adaptive decision rules as realistic features of firm behavior (e.g., Blinder et al., 1998; Camerer, 2003; Zbaracki et al., 2004; Ellison, 2006; Sandholm, 2010; Spiegel, 2014). Alternatively, we can think of the firms as facing significant capital constraints and any forward-looking strategy that is distinct from the short-term strategy must sacrifice current profits for future profits, which the capital constrained firms cannot do.

⁵The results generalize to arbitrary (well-behaved) individual demand structures.

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. Though somewhat obvious (but distinct from adaptive search models such as Lewis, 2011), this dynamic has important implications for the evolution of prices and the potential for convergence.

The short-run is characterized by Edgeworth cycles, a finding unique to the undirected search literature. 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 process then repeats itself.⁶ In the long run, these Edgeworth cycles may persist or the market may converge to the Diamond paradox or a kinked demand equilibrium at the search cost. 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. We show that the upper bound of the cycle is generally not the monopoly price. The lower bound is generally above both the marginal cost and the search cost, rendering each less informative. Moreover, both the upper and lower bounds of the cycles shift with aggregate search behavior (though not necessarily in the same direction), making learning from past experiences difficult.

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 random 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. Because there is residual demand for the firm with the higher price, a firm relents because the residual profit exceeds that of

⁶Edgeworth 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.

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.⁷ Thus, firms do not generally compete the price down to marginal cost.⁸

Nonzero residual demand provides a greater incentive to relent to the firm with the larger installed base of consumers. This result provides a testable implication that larger firms will tend to be the first to raise prices during price wars and 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 of continuing the price war. This linkage 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, and in contrast to the models in the tradition of Varian (1980) or Maskin and Tirole (1988), this model endogenizes how the intensity of competition evolves over time. Rather than a fixed probability of searching or random variance in their search behavior, consumers’ willingness to search responds to the current observed price. This mechanism causes both the peak and trough of the Edgeworth cycle to shift with search activity, as is observed in retail gasoline markets (Noel, 2007a). Empirically, this result implies that measures such as price-comparison traffic, advertising, or price alerts should correlate not only with the troughs of price cycles (as in Lewis and Marvel, 2011), but also with their peaks; however, the correlation with peaks can be positive or negative. Thus, the model provides a micro-foundation for time-varying competition intensity that is absent from canonical search and cycling models. 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

⁷Wallner (1999) finds a deterministic reset driven by cycles being exactly three steps in length, so firms alternate on resets.

⁸Cycling 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.

price stickiness: opportunities for price revisions are themselves random.

The convergence (or lack thereof) of prices and consumer search thresholds depends on the difference between the peak (residual maximizer) and trough (judo price) of the Edgeworth cycle when consumers are least likely to engage in search. If this gap is sufficiently large, neither prices nor the consumers' search thresholds settle and aperiodic stochastic Edgeworth Cycles persist. If, on the other hand, this gap between the judo price and residual maximizer is small, then consumers gradually adopt the highest thresholds and the Edgeworth Cycles persist in a stable manner with fixed bounds. Consequently, if consumers search sufficiently infrequently at the highest threshold, then the cycles shrink, converging to the monopoly price: the Diamond paradox. There is also a special case (when setting the price equal to the search cost maximizes residual profits given the maximum search threshold by all consumers). Instead of converging to the Diamond paradox, prices converge to the search cost. Taking the search cost to zero yields the Bertrand paradox, though through random perturbations (as some consumers will search), prices will rise above zero, so the Bertrand paradox is not asymptotically stable, though it is Lyapunov stable.

The remainder of the paper is structured as follows. Section 2 presents the model. The comparative statics and dynamics are presented in Section 3. We provide an analytical and computational example in Section 4. Discussion and concluding remarks are provided in Section 5. An online appendix provides additional results and expands upon the assumptions made in the model.

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.⁹ With this stochastic continuous time approach, the following two events can occur with

⁹The 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.

positive probability during any open interval of time: the firms update their prices multiple times before consumers update their thresholds or consumers update their thresholds multiple times before the firms update their prices.¹⁰

2.1 Preliminaries and Timing

There are two identical single-product firms competing in prices and a continuum of consumers with unit mass and identical consumption preferences.¹¹ A typical firm is indexed by i . Time flows continuously and is indexed by $t \in [0, \infty)$. Denote by $p^t = (p_1^t, p_2^t)$ the vector of firms' prices at each time t . Let ξ denote 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, where 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 with rate λ_F , selects its price from an ordered finite grid $G = \{g_0, g_1, \dots, g_M\}$ to maximize its instantaneous profits.¹² 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 $\alpha_i \in (0, 1)$.¹³ 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 .

¹⁰Though 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.

¹¹Online Appendix A extends the model to N firms.

¹²For tractability and consistency with observed short-run pricing behavior, firms are assumed to be myopic. That is, they optimize current profits without internalizing how present prices influence the future distribution of consumer search thresholds, a high-dimensional state variable that renders full dynamic optimization infeasible. See, e.g., Weintraub et al. (2008).

¹³For example, suppose that there is a city with two gas stations each located on opposite sides 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. This assumption simplifies the analysis and notation, but does not influence the dynamics or equilibrium outcomes. See Online Appendix A.2 for details.

DATE 3. Each consumer receives opportunities to update her threshold (detailed below) according to an independent Poisson process with rate λ_C . We place no restrictions on the relative magnitudes of λ_C and λ_F insofar as the likelihood of any party making a revision is distinguishable from zero.

The initial price vector p^0 is exogenously fixed, though the equilibrium dynamics do not depend on this starting value. Before proceeding, two remarks are in order.

2.2 Consumer Preferences and Search Strategies

At each instant, each consumer is endowed with unit demand and the stationary utility function $u(\xi) = 1 - \xi$ when purchasing a unit of the good at price ξ and $u(\xi) = 0$ when not purchasing.¹⁴ We make the unit valuation assumption to normalize the monopoly price at 1. The consumers' indirect utility function is $v(\xi, s) = u(\xi) - cs$, where $s = 1$ if the consumer searches and $s = 0$ otherwise.

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 on the ordered grid $\tau \in \{\tau_0, \dots, \tau_L\}$, where $L \geq 1$.¹⁵ These thresholds correspond to the consumers' strategies. Denote by $x^t = (x_0^t, \dots, x_L^t)$ the mass of consumers endowed with each threshold at time t . We assume $x_k^0 > 0$ for all k . Otherwise, any k with $x_k^0 = 0$ can be removed from the grid without any loss of generality.

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 shocked threshold $\tau + \sigma^t$ and the search cost c : $\xi > \max\{\tau + \sigma^t, c\}$.¹⁶ Let $\tau + \sigma^t$ denote the *reservation price*, where any observed price $\xi > \tau + \sigma^t$ leaves that consumer dissatisfied and willing to search.¹⁷ 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

¹⁴All subsequent results hold for both generic unit demands and downward-sloping demands $q = D(\xi)$ such that $\xi D(\xi)$ is strictly quasiconcave.

¹⁵The results generalize to the case of continuous thresholds by augmenting the model following Cheung (2016) and Cheung and Wu (2018).

¹⁶The 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.

¹⁷The 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.

justify the search. The shock σ^t is independent across time and distributed according to the continuous CDF φ . A consumer that observes a price ξ will therefore search with probability $\varphi(\xi - \tau)\iota_{(\xi > c)}$, where $\iota_{(\xi > c)}$ is the indicator function for $\xi > c$.

This reservation price mechanism is distinct from the literature. In this model, reservation prices follow a behavioral rule of thumb: consumers have some level of price sensitivity that evolves over time, but also experience random shocks to their current sensitivity. Search is then governed determined by the observed price compared to this shocked value. In the literature, the reservation price equilibrium is generally derived from consumers forming rational expectations over the distribution of prices (or quality). See, e.g., Varian (1980), Wolinsky (1986), Stahl (1989), Baye et al. (2006), Chen (2024), Janssen and Williams (2024). Alternative mechanisms in the literature include a stochastic minimax stopping rule, where for high prices, consumers randomize between search and not (Parakhonyak and Sobolev, 2015) and rational probabilistic search with Bayesian updating (Janssen et al., 2017).¹⁸

While this reservation price framework yields the same search behavior as models with heterogeneous search costs at any fixed moment in time (a grid of reservation prices with search only if the observed price exceeds the reservation price), the two differ in the dynamics. In the canonical heterogeneous search cost literature (Salop and Stiglitz, 1977; Varian, 1980; Wolinsky, 1986; Stahl, 1989), cross-sectional variation in search costs creates a fixed distribution of reservation prices and stationary mixed-strategy equilibria. Subsequent work generalizes this to continuous heterogeneity in search costs (e.g., Moraga-Gonzalez et al., 2017b,a; Nishida and Remer, 2018). While this creates an observational equivalence to heterogeneous search costs at any moment in time, our thresholds endogenously evolve over time.

We assume that the CDF φ is strictly increasing on $(0 - \tau_L, 1 - \tau_0)$ (increasing the observed price ξ strictly increases the probability of search for any threshold), though the density of this distribution is unrestricted. Thus, the magnitude of the shock may be large enough that a consumer with the highest possible threshold is dissatisfied with any positive price

¹⁸See Anderson and Renault (2018) for a survey.

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.

As the probability of a random consumer with threshold τ_ℓ searching after observing price ξ is $\varphi(\xi - \tau_\ell)\iota_{(\xi > c)}$,

$$\bar{\varphi}(\xi, x) = \sum_{k=0}^L \varphi(\xi - \tau_k) \iota_{(\xi > c)} x_k$$

denotes the probability that a random consumer searches after observing that price. Because there are a continuum of consumers, $\bar{\varphi}(\xi, x)$ is equivalently the mass of consumers that search after observing ξ . Note that $\bar{\varphi}(\xi, x)$ is strictly increasing in ξ on $[c, 1]$, which consistent with the empirical evidence from gasoline markets (Lewis and Marvel, 2011).¹⁹

The distribution of the consumers' thresholds x^t evolves as the consumers update their individual strategies. At each revision opportunity, the consumer employs an imitation dynamic that imposes minimal informational burdens. 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 $r_{k\ell}$. The overall probability that a consumer switches from threshold τ_k to τ_ℓ is thus $\rho_{k\ell} = x_\ell r_{k\ell}$.

Given a pair of strategies, we assume that $r_{k\ell} > r_{\ell k}$ if and only if $v(p^t, s|\tau_k) > v(p^t, s|\tau_\ell)$: 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 is defined by the following system of differential equations for all $k = 0, \dots, L$

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

2.3 Firm Demand, Pricing, and Profits

We now construct each firm's demand as a function of the prices and the distribution of consumers' thresholds. If firm i 's price is lower than its competitor's price, then i serves

¹⁹We prove this and other supporting results in the Appendix.

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

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

We refer to lower priced firm's demand as the *front-side demand* and the higher priced firm's demand as the *residual demand*. The *front-side profit* $\pi_i^F(p, x)$ and *residual profit* $\pi_i^R(\xi, x)$ are defined analogously:

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

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 we show below, the nonnegative residual demand yields distinct equilibrium behavior.

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).²⁰ We then explore how x evolves along the equilibrium

²⁰In Online Appendix C.2, we informally discuss how the results extend to the case of forward-looking firms, where Markov perfection is a more substantial condition.

path. We will show below that first order stochastic dominance offers a natural way to characterize this evolution. In this context, x first order stochastically dominates x' if

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

for all $k = 0, \dots, L$, and strictly so for at least one k .

3.1 Residual Maximizers, Judo Prices, and Best Responses

Define the set of residual maximizers as $\tilde{P}(x) := \arg \max_{g \in G} \pi_i^R(g, x)$, which is independent of i as each firm's residual profit function is a constant multiple of the other. Given the unit demand and G , $\tilde{P}(x)$ is nonempty and any residual maximizer cannot exceed the monopoly price of 1. Thus, all prices are henceforth restricted to be weakly below the monopoly price ($p_i^t \leq 1$) as there is no justification for any firm to price above 1.

The equilibrium characterization is based on the *judo price* of each firm. 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.²¹ Formally,

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

Let $\|G\| := \max_{\omega \geq 1} g_{\omega} - g_{\omega-1}$ denote the norm of G , which we henceforth assume to be sufficiently small. Under this assumption, we show in Lemma 2 (contained in the Appendix) that $p_i^*(x) \in [c, \xi^m]$. We now characterize the firms' best response correspondences.

Proposition 1. *Firm i 's best response correspondence is*

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

²¹The 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.

Proposition 1 demonstrates that a firm will undercut its competitor's price unless that price is below the judo price.²² 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 $\tilde{P}(x)$.²³

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

Proposition 2. *If $\alpha_i > \frac{1}{2}$, then $p_i^*(x) \geq p_{-i}^*(x)$, strictly so if $p_i^*(x) > c$.*

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 competition that the firm with the larger capacity has a higher judo price.²⁴ The following Proposition relates the judo price to the distribution of consumer thresholds.

Proposition 3. *If x first order stochastically dominates x' , then $p_i^*(x) \geq p_i^*(x')$.*

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.

The same relationship need not hold for the residual maximizers. The relationship between $\tilde{P}(x)$ and x depends on the shape of φ . In special cases, $\tilde{P}(x)$ moves with $p_i^*(x)$.²⁵ The residual maximizer determines how the market adjusts as the distribution of search thresholds evolves. When $\tilde{P}(x)$ responds to shifts in x (which occurs when consumers respond to the current observed price), variations in search behavior translate directly into movements

²²In Online Appendix C.2, we show that under certain conditions a similar best response correspondence can be constructed for forward looking firms.

²³This statement is formally proven in Lemma 4 in the Appendix.

²⁴See, e.g., Osborne and Pitchik (1986), Deneckere and Kovenock (1992), and Allison and Lepore (2025).

²⁵Examples offered in Online Appendix B.1 and B.2 demonstrate that the residual maximizer can move in either direction. We formalize conditions under which the residual maximizer moves with the judo price in Online Appendix B.3.

in firms' demand and pricing incentives. Endogenizing consumer thresholds thus provides a mechanism through which observable fluctuations in attention like those driven by advertising intensity, platform traffic, or information shocks, map into changes in prices. This allows the model to explain persistent variation in price dispersion and markups without appealing to exogenous shifts in cost or demand fundamentals.

3.2 Equilibrium Dynamics

Consumers will gradually adopt higher search thresholds when search has a negative average utility and will otherwise gradually adopt lower search thresholds. Define $c^*(p)$ by

$$c^*(p) = \begin{cases} (1 - \alpha)(p_2 - p_1) & \text{if } p_1 < p_2 \\ \alpha(p_1 - p_2) & \text{if } p_2 < p_1. \end{cases}$$

For a given price vector p , $c^*(p)$ represents the largest search cost such that search yields positive expected utility.²⁶

Proposition 4. *Let p remain fixed over any interval of time containing t' and $t > t'$. For all such intervals,*

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

The case in which $c = c^*(p)$ can be ignored as the grid can always be perturbed such that there are no prices that satisfy this relationship. That the distribution of consumer thresholds is always increasing or decreasing (under first order stochastic dominance) is particularly useful because it implies that Proposition 4 characterizes the motion of the firms' judo prices over time (by Proposition 3) and thus the lower bound of the best response (Propositions 1 and 2).

3.2.1 Stochastic Edgeworth Cycles

Equilibrium pricing consists 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

²⁶Formally, this expression is given by $c^*(p) := \sup \{c \geq 0 : v(\min p, 1) > \alpha v(p_1, 0) + (1 - \alpha)v(p_2, 0)\}$.




Figure 1: Edgeworth cycles for a fixed distribution of thresholds x . For simplicity, we assume a unique residual maximizer for the purposes of the figure.

the constantly changing consumer thresholds and stochastic nature of price stickiness, the actual pattern of pricing is not cyclical in the classical sense: the bounds of the price war are not constant. Nonetheless, the general pattern repeats and maintains the fundamental characteristics of an Edgeworth cycle.

Define $p^*(x) = \max \{p_1^*(x), p_2^*(x)\}$ as the *critical judo price*. 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 ($\alpha \leq 1/2$), and thus by Proposition 2 the higher judo price. Hence, firm 2's judo price is also the critical judo price. The Edgeworth cycle proceeds as follows, with an example depicted in Figure 1:

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 $p_2^*(x)$.
4. Firm 2 relents and sets a price in $\tilde{P}(x)$ to maximize its residual profit.
5. Repeat this process from step 1.

During steps 1-3 of the cycle, the prices of the firms will be close enough that search will not

be beneficial. Thus, Proposition 4 implies that consumers will be adopting higher search thresholds, which by Proposition 3 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 2 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. As consumers raise their thresholds and the judo prices increase, the critical judo price may increase to a value above the current price. 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.

While appearing as a mixed strategy due to the stochastic timing of consumer behavior, relenting (step 4) and resetting the Edgeworth cycle is a pure strategy. This pattern is in contrast to typical models of Edgeworth cycles such as Maskin and Tirole (1988), Wallner (1999), and Noel (2008). Differences in the relenting firm can be attributed to and empirically identified as changes in consumer search behavior.

The equilibrium path and convergence (or lack thereof) depends on the parameters and functional forms of the model, namely the search cost and distribution φ . This section considers two conditions and the subsequent section analyzes their complements. Denote by e_k 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 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 (via φ) in that it requires $p^*(e_L) < \inf \tilde{P}(e_L)$.

The following Theorem demonstrates that the equilibrium pricing dynamic does not converge under this condition and that the range of prices in the cycles is large enough to induce search. Define $\underline{p} = p^*(e_0)$ as the critical judo price when all consumers have the

lowest search threshold and $\hat{p} = \inf \tilde{P}(e_L)$ as the smallest residual maximizer when all consumers have the highest search threshold.

Theorem 1. *Under C1, there exists a $T > 0$ such that:*

- (i) *neither p^t nor x^t converge as $t \rightarrow \infty$,*
- (ii) *for all times $t > T$, $p_i^t \in [\underline{p}, 1]$,*
- (iii) *$(x^t, p_i^t) \rightarrow (e_0, \underline{p})$ and $(x^t, p_i^t) \rightarrow (e_L, \hat{p})$ 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 e_0 and e_L . Proposition 1 implies that a firm will eventually choose a price that is just below its competitor's price. When the grid is sufficiently fine, the difference in prices is less than the cost of search so it cannot be beneficial to search. By proposition 4, 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 .²⁷ 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 $p^*(e_L)$ 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 C2, there exists a $T > 0$ and an equilibrium such that*

²⁷It merits mention that if $\lambda_F - \lambda_C \gg 0$, then the probability of the consumer dynamic moving from the two extremes becomes arbitrarily small.

- (i) neither p^t nor x^t converge as $t \rightarrow \infty$,
- (ii) for all times $t > T$, $p_i^t \in [\underline{p}, 1]$, and
- (iii) $(x^t, p_i^t) \rightarrow (e_0, \underline{p})$ and $(x^t, p_i^t) \rightarrow (e_L, \hat{p})$ infinitely many times.

C1 and C2 coincide if and only if $\tilde{P}(e_L)$ is a singleton.²⁸ 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 of search costs if $\tilde{P}(e_L)$ is not a singleton, though the pricing strategy is unique (Proposition 1).

The mechanism driving Theorem and Corollary 1 is fundamentally different than Maskin and Tirole (1988) and others found in the literature. In Maskin and Tirole’s seminal work, cycles arise from forward-looking best responses and mixed-strategy relenting among firms. In contrast, the cycles here originate from demand-side bounded rationality: consumers’ evolving search thresholds create positive residual demand for the higher-priced firm, making relenting a deterministic and profitable action that endogenously determines both its timing and identity. The potential existence of multiple residual maximizers is also important, as it produces distinct equilibrium paths and variation in cycle amplitude and persistence that could vanish under a uniqueness assumption. Moreover, unlike earlier stationary-cycle frameworks, whether operating under capacity-constraints (Edgeworth, 1925; Shubik, 1959; Osborne and Pitchik, 1986) or the dynamic mixed-strategy models of Maskin and Tirole (1988), Wallner (1999), and Eckert (2003), the upper and lower bounds of the cycle in our model move endogenously with consumer search behavior. This non-stationarity mirrors empirical evidence from retail gasoline markets, where price-cycle peaks and troughs shift independently of cost movements (e.g. Noel, 2007a,b; Eckert and West, 2004; Wang, 2009; Doyle et al., 2010; Zimmerman et al., 2012; Isakower and Wang, 2014).

²⁸The arguments made to prove this corollary are nearly identical to those made in Theorem 1 with one key difference: G must be sufficiently fine such that for some neighborhood $\mathcal{N}(e_L)$ of e_L and all $x \in \mathcal{N}(e_L)$, $\sup \tilde{P}(e_L) \in \tilde{P}(x)$ and the equilibrium needs to dictate that when a firm chooses a price $p^t \in \tilde{P}(x)$, it will choose $p^t = \sup \tilde{P}(e_L)$.

3.2.2 Limit Cycles, 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 to a fixed peak and trough (though aperiodicity persists due to price stickiness). 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'.²⁹ 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 C2', there exists a $T > 0$ such that*

- (i) *the distribution of consumer thresholds x^t converges to e_L as $t \rightarrow \infty$,*
- (ii) *for all times $t > T$, $p_i^t \in [p^*(e_L), \sup \tilde{P}(e_L)]$, and*
- (iii) *$p^t \rightarrow p^*(e_L)$ and $p_i^t \rightarrow \xi \geq \hat{p}$ infinitely many times.*

²⁹To 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'.

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 never 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.

Corollary 2. *Under C1', there exists a $T > 0$ and an equilibrium such that*

- (i) *the distribution of consumer thresholds x^t converges to e_L as $t \rightarrow \infty$,*
- (ii) *for all times $t > T$, $p_i^t \in [p^*(e_L), \sup \tilde{P}(e_L)]$, and*
- (iii) *$p^t \rightarrow p^*(e_L)$ and $p_i^t \rightarrow \xi \geq \hat{p}$ 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 5. *If τ_L is sufficiently large, then there is convergence to the Diamond Paradox ($p_i^t \rightarrow 1$ as $t \rightarrow \infty$).*

This result is similar to the Diamond paradox that emerges in the dynamic model of undirected consumer search in Diamond (1971) and the static model of Stahl (1989). The difference lies in the mechanism. In this model, firms price so as to gradually discourage

search, and as consumers search sufficiently infrequently, monopoly pricing becomes optimal. Once both firms reach this point, search ceases to be beneficial and the outcome stabilizes. In contrast, the Diamond-Stahl mechanism relies on consumers forming rational expectations over equilibrium price distributions; monopoly pricing emerges because both sides anticipate that outcome.

Proposition 5 generalizes this logic by linking the Diamond-type outcome to the kinked-demand equilibrium of Maskin and Tirole (1988). In Stahl (1989), convergence to monopoly pricing and no search occurs when the exogenously determined proportion of searchers tends to zero. In Maskin and Tirole (1988), convergence to a kinked-demand equilibrium follows from the supply side: infinitely patient firms sustain monopoly pricing as a Markov perfect equilibrium by internalizing the future cost of price wars. In contrast, convergence here arises entirely from demand-side dynamics. Consumers' search behavior is endogenously driven by the firms' pricing: as prices remain close, consumers reduce search intensity, raising the lower bound of the cycles (Proposition 3). Under Condition 2, these adjustments are insufficient to reverse the process, and prices gradually approach the monopoly level as the distribution of thresholds converges to 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 C3, there exists a $T > 0$ such that $x^t \rightarrow e_L$ 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 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 Analytical Example

This section provides an analytical example and computational simulation to illustrate Theorems 1 and 2 and Proposition 5. For this section, suppose that the distribution of shocks σ^t is uniformly distributed between -1 and 1 . Hence,

$$\varphi(\xi, \tau) = \frac{1 + \xi - \tau}{2}.$$

For expositional convenience, define $\bar{\tau}(x) = \sum_{\ell=1}^L x_\ell \tau_\ell$, so

$$\bar{\varphi}(\xi, x) = \frac{1 + \xi - \bar{\tau}(x)}{2}.$$

Then,

$$\begin{aligned} \arg \max_{\xi} \underbrace{\alpha \xi (1 - \bar{\varphi}(\xi, x))}_{\pi_R(\xi, x)} &= \arg \max_{\xi \in G} \xi (1 + \xi - \bar{\tau}(x)) \\ &= \frac{1 + \bar{\tau}(x)}{2}. \end{aligned} \tag{1}$$

As the residual profits are strictly concave, $\tilde{P}(x)$ is unique and given by the $g \in G$ nearest to (1). Let $\tilde{g}(x) \in G \cap \tilde{P}(x)$ denote the residual maximizer. Hence,

$$D_i(p, x) = \begin{cases} \alpha_i + (1 - \alpha_i) \left(\frac{1 + p_{-i} - \bar{\tau}(x)}{2} \right) & \text{if } p_i < p_{-i} \\ \alpha_i & \text{if } p_i = p_{-i} \\ \alpha_i \left(\frac{1 - p_i + \bar{\tau}(x)}{2} \right) & \text{if } p_i > p_{-i}. \end{cases}$$

We can then specify the judo price as the largest g_ω that satisfies the minimum of the search cost c and

$$g_{\omega-1} \left(\alpha_i + (1 - \alpha_i) \left(\frac{1 + g_\omega - \bar{\tau}(x)}{2} \right) \right) \leq \alpha_i \tilde{g}(x) \left(\frac{1 - \tilde{g}(x) + \bar{\tau}(x)}{2} \right).$$

Suppose that $g_\omega - g_{\omega-1} = \delta$ for all $\omega \geq 1$ and G is sufficiently fine (δ is sufficiently small). Then $g_{\omega-1} = g_\omega - \delta$ and $|\tilde{g}(x) - \frac{1+\bar{\tau}(x)}{2}| < \varepsilon$ for an arbitrarily small ε . Then, the judo price is (approximately) given by the minimum of the search cost c and the largest g_ω that satisfies

$$(g_\omega - \delta) \left(\alpha_i + (1 - \alpha_i) \left(\frac{1 + g_\omega - \bar{\tau}(x)}{2} \right) \right) = \frac{\alpha_i}{8} (1 + \bar{\tau}(x))^2,$$

which can be rewritten as

$$\begin{aligned} \left(\frac{1 - \alpha_i}{2} \right) p^*(x)^2 + \left(\alpha_i + \left(\frac{1 - \alpha_i}{2} \right) (1 - \bar{\tau}(x) - \delta) \right) p^*(x) \\ = \delta \left(\alpha_i + \left(\frac{1 - \alpha_i}{2} \right) (1 - \bar{\tau}(x)) \right) + \frac{\alpha_i}{8} (1 + \bar{\tau}(x))^2. \end{aligned}$$


A standard implicit function theorem argument confirms Propositions 2 and 3.

Observe that as $x \rightarrow e_L$, $\tilde{g}(x) \rightarrow 1$ and as $x \rightarrow e_0$, $\tilde{g}(x) \rightarrow \frac{1}{2}$. Thus the peak of the cycles varies between $\frac{1}{2}$ and 1 and moves uniformly with x (with respect to first order stochastic dominance). Similarly, the troughs of the cycles move with x uniformly (Proposition 3).


4.1 Computational Simulation

In this section, we conduct a computational simulation to visualize the main results of Theorem 1 and Proposition 5. To conduct this simulation, we make a few revisions to the model; particularly moving from continuous time to discrete time and from a continuum of consumers to a finite population; however, these revisions (described below), still keep the framework consistent with our setup in Section 2.


Suppose that $G = \{0, 0.01, 0.02, \dots, 0.99, 1\}$ ($\delta = 0.01$), and $\alpha = 0.6$. The number of consumers is set to 50. The simulation operates in discrete time as follows. In each period, either a single consumer or single firm has the opportunity to revise their strategy. With probability 0.5, a random consumer is drawn to update its search threshold and with probability 0.5, a random firm is given the opportunity to update its price. In the event a firm i is revising, the firm best responds to the current state x and p_{-i} according to Proposition 1. Consumer thresholds remain unchanged. In the event a consumer i is revising, they observe firm 1's price with probability $\alpha = 0.6$ and firm 2's price with complementary probability. The consumer receives a stochastic shock to their threshold given by the uniform



(a) $\tau = \{0, 1.5\}$ and $c = 0.02$.



(b) $\tau = \{0, 1.5\}$ and $c = 0.05$.



(c) $\tau = \{0, 2\}$ and $c = 0.02$.

Figure 2: The short- and long-run dynamics. Panel (a) satisfies condition $C1$ (Theorem 1), Panel (b) satisfies condition $C2'$ (Theorem 2), and Panel (c) satisfies the conditions of Proposition 5. The gray lines correspond to prices and the black lines correspond to consumer thresholds.

distribution described above, and decide whether to search or not. A purchase decision is then made. After purchasing, with probability 0.001, consumer i is replaced with a new consumer with a random threshold, which simulates the continuum of consumers. With complementary probability, consumer i is matched with a random consumer j , and with probability $v_i/(v_i + v_j)$, adopts consumer j 's threshold, where v_i v_j are the consumers' respective indirect utility functions. The process then repeats in the next period.

Figure 2 illustrates the dynamics under two conditions: (a) $\tau = \{0, 1.5\}$ and $c = 0.02$, (b) $\tau = \{0, 1.5\}$ and $c = 0.05$, and (c) $\tau = \{0, 2\}$ and $c = 0.02$. Under the first condition, C1 is satisfied so the nonconvergence results of Theorem 1 emerge. There are persistent cycles that fluctuate. Both the peak and trough of the cycle move together with x with respect to first order stochastic dominance. Under (b), Condition C2' is satisfied, yielding persistent cycles in the limit with $x \rightarrow e_L$. Under (c), the highest threshold is enough that the probability of search is zero, yielding the Diamond Paradox, though cycles persist in the short run leading to the long run convergence at the monopoly price. The judo price gradually shifts upward until both the judo price and residual maximizer converge to the monopoly price.

5 Discussion and Concluding Remarks

By developing a model of undirected consumer search with firm competition, this paper characterizes the resulting market outcomes and derives several empirically testable predictions, including the short- and long-run dynamics that arise from agents' bounded rationality. The short-run dynamics take the form of stochastic Edgeworth cycles, which largely persist in the long run except under specific conditions. For a given distribution of consumer search thresholds, the firm with the larger installed base has a stronger incentive to monopolize residual demand rather than continue the price war. Consequently, when firms can update prices more frequently than consumers can reevaluate their search decisions, larger firms are more likely to end and reset the cycle by monopolizing residual demand. This prediction is consistent with evidence from gasoline markets (Noel, 2007b; Atkinson, 2009; Zimmerman et al., 2012; Isakower and Wang, 2014). Conditions C1, C2', and C3 determine long-run convergence. Although expressed in terms of the judo price, residual maximizer, and search cost, these conditions can be traced back to the model's fundamentals, and convergence (or its absence) depends on exogenous factors.

The model offers a unified explanation for the distinct pricing regimes documented in Noel (2007a), who identifies three patterns in retail gasoline markets: sticky pricing, Edgeworth price cycles, and cost-based pricing. Each of these regimes corresponds to a specific theoretical result in this paper: Edgeworth cycles to Theorems 1 and 2 and cost-based and sticky

pricing to Theorem 3 and Proposition 5. Rather than relying on equilibrium selection, the model explains observed variation across cities by potentially subtle differences in consumer demand and search costs. A formal empirical analysis of this mechanism is left for future research.

The bounded rationality of consumers has important implications. The model does not coincide with Maskin and Tirole (1988), Eckert (2003), or Noel (2008), even if search costs are zero or firms are forward-looking. Whenever prices equalize, consumers randomly shift toward higher price thresholds, creating an incentive for firms to raise prices. Both equilibria characterized by Maskin and Tirole involve firms setting identical prices for at least some period. Thus, the underlying consumer dynamics in this model would induce a deviation from their results that does not vanish in the limit.

While the firms are modeled as myopic, this assumption can be both behaviorally and theoretically justified. The analysis focuses on short-run price dynamics rather than long-run equilibrium refinement. In this environment, the firms' behavior can be interpreted as an adaptive best response to current profit opportunities. Forward-looking optimization would require firms to internalize how their pricing decisions affect the evolution of the entire distribution of consumer search thresholds, a high-dimensional state variable that makes a fully dynamic programming approach analytically infeasible. Myopia isolates the demand-side feedback mechanism driving the Edgeworth cycles. The assumption is also consistent with the broader literature on bounded rationality in industrial organization, which emphasizes the empirical relevance of limited foresight and adaptive decision rules in competitive environments (e.g., Camerer, 2003; Blinder et al., 1998; Zbaracki et al., 2004; Ellison, 2006; Sandholm, 2010; Spiegler, 2014). Empirically, such short-horizon pricing behavior aligns with evidence from managerial pricing studies (Blinder et al., 1998; Zbaracki et al., 2004), which show that firms frequently rely on adaptive rather than fully rational forward-looking strategies. However, as we illustrate in Online Appendix C, the results are not driven by this assumption.

A major implication of the model is that the Bertrand paradox does not emerge as an equilibrium outcome when the cost of search tends to zero. This result demonstrates the

asymptotic instability of the Bertrand outcome, which occurs only under highly specific assumptions and is overturned by infinitesimal perturbations. Marginal-cost pricing cannot persist even with arbitrarily small search costs, since consumers have no incentive to search when prices approach marginal cost. Over time, random shifts in the distribution of price thresholds (Proposition 4) allow firms to earn positive profits. Hence, the Bertrand outcome is only a transient feature of the degenerate case with exactly zero search costs.

The level of the search cost does not affect the qualitative nature of the results as long as it remains positive and the pricing grid is sufficiently fine. Search cost matters only insofar as it determines whether search has positive or negative expected value. Its primary influence is on the rate at which the process evolves. Because the probability that a consumer adopts a strategy increases with the expected payoff from that strategy, a higher search cost reduces the expected value of search in all states. This accelerates consumers' increases in thresholds during cycles and slows their lowering when the cycle resets.

Although the magnitude of search cost does not alter the qualitative results, the model offers insight into its role in pricing incentives. Higher search costs accelerate convergence to higher price thresholds and slow convergence to lower ones, so firms benefit when consumers search less frequently. Moreover, both the critical judo price and the smallest residual maximizer are bounded below by the search cost. Firms can therefore increase profits by raising the effective cost of search through obfuscation (Ellison and Ellison, 2009; Ellison and Wolitsky, 2012). When binding, such as when the judo price equals the search cost or under Condition 3 and the assumptions of Theorem 3, firms can strictly benefit from actions that increase c , provided these actions are not excessively costly.

Finally, most models of undirected search include a fraction of “shoppers” consumers who either face zero search cost or are exogenously informed of firms' prices. Because this paper's goal is to weaken the standard rationality assumptions and study boundedly rational dynamics, such an inclusion would be inappropriate here. Nevertheless, the implications are straightforward. Shoppers increase the share of consumers who search given any observed price. Their inclusion is equivalent to introducing a negative bias in each consumer's price threshold. If the fraction of shoppers is sufficiently large, the conditions of Theorem 1

hold, leading to perpetual cycling of prices and thresholds. If the fraction is small, the conditions of Theorem 2 and Proposition 5 may still apply, yielding long-run convergence to the Diamond paradox.

This result contrasts with Stahl (1989), who shows in a static rational-expectations model that equilibrium prices vary continuously between the Bertrand and Diamond outcomes as the fraction of shoppers moves from zero to one.³⁰ Similarly, it differs from the theoretical results of Pennerstorfer et al. (2020), who obtain comparable predictions to Stahl (1989). However, the persistence of price dispersion in this model is consistent with the semiparametric empirical evidence in Pennerstorfer et al. (2020), where the degree of price dispersion declines but does not vanish as the share of shoppers approaches zero or one.

References

- Allison, Blake A. and Jason J. Lepore**, “A general model of Bertrand Edgeworth duopoly,” *Games*, 2025, 16 (3), 26.
- 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.
- and —, “Chapter 8: Firm Pricing with Consumer Search,” in Luis C. Corchón and Marco A. Marini, eds., *Handbook of Game Theory and Industrial Organization*, Vol. 2 2018.
- 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-specific 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.
- Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd**, *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation, 1998.

³⁰Equilibrium pricing in Stahl’s model is in mixed strategies when the fraction of shoppers lies strictly between zero and one.

- 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.
- Camerer, Colin F.**, *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, 2003.
- Castanias, Rick and Herb Johnson**, “Gas wars: Retail gasoline fluctuations,” *Review of Economics and Statistics*, 1993, *75* (1), 171–174.
- Chen, Yongmin**, “Search and competition under product quality uncertainty,” *Journal of Industrial Economics*, 2024, *72* (2), 633–661.
- 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.
- Conlisk, John**, “Why bounded rationality?,” *Journal of Economic Literature*, 1996, *34* (2), 669–700.
- 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.
- Heidhues, Paul and Botond Köszegi**, “Behavioral Industrial Organization,” in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics – Foundations and Applications 1*, Amsterdam: North-Holland (Elsevier), 2018.
- 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.
- **and Cole Williams**, “Consumer Search and Product Returns in E-Commerce,” *American Economic Journal: Microeconomics*, 2024, 16 (2), 387–419.
- **, 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.
- Moraga-Gonzalez, José Luis, Zsolt Sándor, and Matthijs R. Wildenbeest**, “Nonsequential search equilibrium with search cost heterogeneity,” *International Journal of Industrial Organization*, 2017, 50, 392–414.
- **, —, and —**, “Prices and heterogeneous search costs,” *RAND Journal of Economics*, 2017, 48 (1), 125–146.
- Nishida, Mitsukuni and Marc Remer**, “Lowering consumer search costs can lead to higher prices,” *Economics Letters*, 2018, 162, 1–4.
- 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.
- Parakhonyak, Alexei and Anton Sobolev**, “Non-reservation price equilibrium and search without priors,” *The Economic Journal*, 2015, 125 (584), 887–909.

- 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: Massachusetts 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 University 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**, “Sequenal 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.
- Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy**, “Markov perfect industry dynamics with many firms: A curse of dimensionality,” *Econometrica*, 2008, 76 (6), 1375–1411.
- Wolinsky, Asher**, “True monopolistic competition as a result of imperfect information,” *Quarterly Journal of Economics*, 1986, 101 (3), 493–512.
- Yang, Huanxing and Lixin Ye**, “Search with learning: Understanding asymmetric price adjustments,” *RAND Journal of Economics*, 2008, 39 (2), 547–564.
- Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen**, “Managerial and customer costs of price adjustment: Direct evidence from industrial markets,” *Review of Economics and Statistics*, 2004, 119 (2), 1125–1160.

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. *Define X as the unit simplex in $\mathbb{R}^L$:*

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

. 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 the assumption on φ 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

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

which simplifies to

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

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

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

which is true if x exhibits first order stochastic dominance over x' . $\square$

Lemma 2. Suppose $\|G\| < \delta$ for some $\delta > 0$. Then, $c \leq p_i^*(x) < 1$.

Proof of Lemma 2.

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

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

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 the properties of φ , $\bar{\varphi}(1, x) > 0$. Consider $\xi \in \tilde{P}(x)$ and note that

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

Hence, $p_i^*(x) < 1$. $\square$

The proof of Proposition 1 relies on the following three lemmas.

Lemma 3. Let $g_{\omega^*(x)} = p_i^*(x)$. There exists a $\delta > 0$ such that if $\|G\| < \delta$, then $g_{\omega^*(x)} < g_M = 1$.

Proof. Suppose to the contrary that $g_{\omega^*(x)} = g_M = 1$ and consider $\xi \in \tilde{P}(x)$. The proof proceeds in two cases: $\xi < 1$ and $\xi = 1$.

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

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

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

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

By (2) and (3), $\alpha_i \xi(1 - \bar{\varphi}(\xi, x)) < \alpha_i \xi(1 - \bar{\varphi}(\xi, x))$, a contradiction.

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

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

which when rearranged yields

$$1 - g_{M-1} > g_{M-1} \frac{\bar{\varphi}(1, x)}{\alpha(1 - \bar{\varphi}(1, x))}. \quad (4)$$

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

$$\delta > (1 - \delta) \frac{\bar{\varphi}(1, x)}{\alpha(1 - \bar{\varphi}(1, x))}. \quad (5)$$

The LHS of (5) is decreasing in δ while the RHS of (5) is increasing in δ . Taking $\delta \rightarrow 1$ yields $1 > 0$ and taking $\delta \rightarrow 0$ yields

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

a contradiction. By continuity, there exists a δ' such that there is a contradiction for all $\delta \leq \delta'$, so for $\|G\| < \delta'$, $g_{\omega^*(x)} = p_i^*(x) < g_M = 1$. $\square$

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

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

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

and

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

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

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

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

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

$$\bar{\varphi}(g_\omega, x) \leq \frac{\alpha_i}{1 - \alpha_i} (g_\omega - g_{\omega-1}) \iff \delta \geq g_\omega - g_{\omega-1} \geq \frac{1 - \alpha_i}{\alpha_i} \bar{\varphi}(g_\omega, x). \quad (7)$$

Choose δ such that

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

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. $\square$

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

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

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

By Lemma 4, $g_\omega \in R_i(g_\omega, x)$ implies that $g_\omega \in \tilde{P}(x)$, so $g_{\omega-1} \in R_i(g_\omega, x)$. $\square$

Proof of Proposition 1.

Proof. Suppose that $\|G\| < \delta < c$. Then $g_0(1 - \bar{\varphi}(g_1, x)) = 0$, while there is some $g \in G$ with $g \in (0, c)$ guaranteeing that

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

By Lemmas 3 and 5,

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

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

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

Proof of Proposition 2.

Proof. As each $p_i^*(x) \geq c$, the result holds trivially if $p_{-i}^*(x) = c$. Suppose that $p_i^*(x) \in (c, 1)$. By the continuity of $\bar{\varphi}(\xi, x)$ on $(c, \infty) \times X$, $\pi_{-i}^F(p_{-i}^*(x), x) = \pi_{-i}^R(x)$. That is,

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

Define the function

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

so that $p_{-i}^*(x)$ is defined (for sufficiently small $\|G\|$ and possibly by a small perturbation of G) 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 :

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

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

Thus

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

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, $p_i^*(x) > p_{-i}^*(x)$. $\square$

Proof of Proposition 3.

Proof. Suppose that $\|G\| < \delta$, where $\delta > 0$ is sufficiently small such that each firm's best response correspondence is as in Proposition 1 and 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, it follows that for all $g_{\omega} \in G$,

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

and

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

If

$$g_{\omega-1}(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_{\omega}, x')) \leq \max_{p_i \in G} \alpha_i p_i (1 - \bar{\varphi}(p_i, x')), \quad (8)$$

then

$$g_{\omega-1}(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_{\omega}, x)) \leq \max_{p_i \in G} \alpha_i p_i (1 - \bar{\varphi}(p_i, x)). \quad (9)$$

Therefore

$$\begin{aligned} & \max \left\{ g_{\omega} \in G \setminus \{g_0\} : g_{\omega-1}(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_{\omega}, x')) < \max_{p_i} \alpha_i p_i (1 - \bar{\varphi}(p_i, x')) \right\} \\ & \leq \max \left\{ g_{\omega} \in G \setminus \{g_0\} : g_{\omega-1}(\alpha_i + (1 - \alpha_i)\bar{\varphi}(g_{\omega}, x)) < \max_{p_i} \alpha_i p_i (1 - \bar{\varphi}(p_i, x)) \right\}, \end{aligned}$$

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

Thus, $p_i^*(x) \geq p_i^*(x')$. Furthermore, the same argument is valid given strict inequalities in (8) and (9). $\square$

Proof of Proposition 4.

Here, we prove the formal statement: suppose that $p^t = p$ for all $t \in [T, T + \varepsilon)$ for any T and $\varepsilon > 0$. 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 ,
- (ii) if $c > c^*(p)$, then x^t first order stochastically dominates $x^{t'}$.

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 .

Suppose $c > c^*(p)$. Then $r_{k\ell} > r_{\ell k}$ if and only if $k < \ell$ because search is not profitable when $c > c^*(p)$. Recall that the net flow of x^t is

$$\begin{aligned}\dot{x}_k^t &= \sum_{\ell} x_{\ell}^t \rho_{\ell k} - x_k^t \sum_{\ell} \rho_{k\ell} \\ &= x_k^t \sum_{\ell} x_{\ell}^t (r_{\ell k} - r_{k\ell}).\end{aligned}\tag{10}$$

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

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

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

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

For all $a = 0 : L$ and by (10),

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

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

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

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

The following Lemma is used to prove Theorem 1.

Lemma 6. *Under 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 $p_i = \inf \tilde{P}(x)$ and $p_{-i} = p_i^*(x) = p^*(x)$ 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 $\tilde{P}(x) = \tilde{P}(e_L)$ for all $x \in \mathcal{N}(e_L)$. Take $g \in \tilde{P}(e_L)$ and define

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

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

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

for all $g' \in G \setminus \tilde{P}(e_L)$ 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 $g' \in G \setminus \tilde{P}(e_L)$. Hence, $g \in \tilde{P}(x)$ and $\tilde{P}(x) = \tilde{P}(e_L)$. As e_L first order stochastically dominates all $x \in X$, $p^*(x) \leq p^*(e_L)$ for all $x \in X$ by Proposition 3.

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

$$1 - \min p - c^*(p) = 1 - \alpha p_1 - (1 - \alpha)p_2. \quad (11)$$

Suppose that $p_1 \leq p_2$. Then, (11) can be rearranged as

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

which by inspection is decreasing in $p_1 = \min p$ and increasing in $p_2 = \max p$. Thus, (11) is decreasing in $\min p$ and increasing in $\max p$.

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 1, δ' can be chosen such that if $\|G\| < \delta'$, then $|p^*(\underline{x}) - p^*(e_L)| \leq \delta(\varepsilon)$. Because $c^*(p)$ is decreasing in $\min p$, a well-defined neighborhood $\mathcal{N}(e_L)$ exists for every ε' such that

$$\begin{aligned} & \min\{c^*(p^*(\underline{x}), \min \tilde{P}(\underline{x})), c^*(\min \tilde{P}(\underline{x}), p^*(\underline{x}))\} \\ & > \min\{c^*((p^*(e_L), \inf \tilde{P}(e_L))), c^*((\inf \tilde{P}(e_L), p^*(e_L)))\} - \varepsilon. \end{aligned} \quad (12)$$

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 $c < \min\{c^*(p^*(x), \min \tilde{P}(x)), c^*(\min \tilde{P}(x), p^*(x))\}$ for all $x \in \mathcal{N}(e_L)$ and G such that $\|G\| < \delta'$. $\square$

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 1.

First, we prove statement (i) (the prices do not converge). To the contrary, suppose that $p^t \rightarrow (g_\omega, g_{\omega'})$. Without loss of generality, assume that $g_\omega \leq g_{\omega'}$. By Proposition 1, for some time $T > 0$, it must that (i) $g_\omega \leq p_2^*(x^t)$ for all $t > T$, (ii) $g_{\omega'} \in \tilde{P}(x)$, 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 4 implies that $x^t \rightarrow e_L$ as $t \rightarrow \infty$ because at such prices, $\dot{x}_L^t \rightarrow 0$ if and only if $x_L \rightarrow 1$. Consider two cases.

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

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

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

Case 2: $g_\omega = g_{\omega'-1}$. In this case, it must be that $g_{\omega'} \geq p_1^*(x^t)$ 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,

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

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 1, given a distribution x^t , neither firm i will ever choose a price $p_i^t < g_{\omega-1}$ when $g_\omega = p_i^*(x)$. Thus, by Propositions 1 and 3, neither firm i will ever choose a price below $p_i^t < g_{\omega-1}$ when $g_\omega = p_i^*(e_0)$. Set $g_\omega = p_i^*(e_0)$ and suppose that $p_i^*(e_0) = p^*(e_0)$. 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 1 guarantees that for sufficiently small $\delta > 0$, if $\|G\| < \delta$, then $p_i^t \geq \underline{p}$ for both firms i and all $t > T$. That $p_i^t \leq 1$ for all $t > T$ follows directly.

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) $p_i^{t'_n} \rightarrow \xi' \geq \hat{p} - \varepsilon$ for some $t > T$, (c) $x^t \in \mathcal{N}(e_0)$ for some $t > T$, and (d) $p_i^{t'_n} \rightarrow \xi \leq \underline{p} + \varepsilon$ for some $t > T$. We will jointly demonstrate (a) and (b) followed by (c) and (d).

Proposition 1 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 4 implies that the distribution of consumer thresholds will be shifting toward e_L .

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

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 $x \in \mathcal{N}(e_L)$, $p_i^*(x) = p_i^*(e_L)$ for each firm i and $\tilde{P}(x) \subseteq \tilde{P}(e_L)$ (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 $p_i \in \tilde{P}(x)$ 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 $(T, T + \tilde{T}(x^T))$ and there is a positive probability that for either firm i , for some time $t > T$, $p_i^t \geq \hat{p}$.

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 $p_i = \inf \tilde{P}(x)$ and $p_{-i} = p_i^*(x) = p^*(x)$. 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 $x \in \mathcal{N}(e_0)$, $p^*(x) = p^*(e_0)$. 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 $p_i = g_\omega = p^*(x)$ 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 $p_i^t \rightarrow p^*(e_L)$ for each firm i for some $t > T$. $\square$

The following Lemma is used to prove Theorem 2.

Lemma 7. *Under 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 $p_i = \sup \tilde{P}(x)$ and $p_{-i} = p_i^*(x) = p^*(x)$.*

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 $\tilde{P}(x) = \tilde{P}(e_L)$. As in the proof of Lemma 6, e_L first order stochastically dominates all $x \in X$, so $p^*(x) \leq p^*(e_L)$ for all $x \in X$.

We now prove that, for all $x \in \mathcal{N}(e_L)$, $c > c^*(p)$ when $p_i = \max \tilde{P}(x)$ and $p_{-i} = p_i^*(x) = p^*(x)$. 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 1, δ' can be chosen such that if $\|G\| < \delta'$, then $\max \tilde{P}(e_L) \geq \sup \tilde{P}(e_L) - \delta(\varepsilon)$. Because $c^*(p)$ is decreasing in $\min p$ and thus $p^*(x)$ while $\max \tilde{P}(x) = \max \tilde{P}(e_L)$ for all $x \in \mathcal{N}(e_L)$,

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

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 $c > \min\{c^*(p^*(x), \max \tilde{P}(x)), c^*(\max \tilde{P}(x), p^*(x))\}$ for all $x \in \mathcal{N}(e_L)$ and G such that $\|G\| < \delta'$. $\square$

Proof of Theorem 2. We prove the following formalized statement of Proposition 5:

As $\tau_L \rightarrow \xi^m - \inf \text{supp } \varphi$, $p_i^*(e_L) \rightarrow \xi^m$ 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 $t > T$.

Proof. Suppose that $\|G\| < \delta < c$, where δ is such that the best response correspondences are as stated in Proposition 1. We first show that $x^t \rightarrow e_L$. By Proposition 1, 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 4, 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 4 imply that e_L is an absorbing state. Hence, $x^t \rightarrow e_L$.

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 1. $\square$

Proof of Proposition 5.

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

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

Hence, $\tilde{P}(e_L) = \{1\}$. Therefore, the judo price is

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

Evaluating firm i 's front-side profits as $p_i \rightarrow 1$ and $p_{-i} = 1$ yields

$$\alpha_i = (\alpha_i + (1 - \alpha_i) \varphi(1 - \tau_L)).$$

Hence, $p_i^*(e_L) = 1 = \sup \tilde{P}(e_L)$. 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 1. By continuity, for every $\varepsilon' > 0$, There exists a neighborhood $\mathcal{N}(1 - \inf \text{supp } \varphi)$ such that if $\tau_L \in \mathcal{N}(1 - \inf \text{supp } \varphi)$, then $p_i^*(e_L) = 1$. Let $\bar{\tau} = \inf \mathcal{N}(1 - \inf \text{supp } \varphi)$. Then, for all $\tau > \bar{\tau}$, $p_i^*(e_L) = 1$.

As C2' is satisfied, Theorem 2 implies that $x^t \rightarrow e_L$. Hence, there exists a time $T \geq 0$ such that for all $t > T$, $p_i^t = 1$. $\square$

Proof of Theorem 3.

Proof. Suppose C3 is satisfied and set δ sufficiently small such that the best response correspondences are as in Proposition 1. 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 1 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 $t > T'$. $\square$