Pricing with algorithms[*]

Rohit Lamba Sergey Zhuk

June 2022

Abstract

This paper studies Markov perfect equilibria in a repeated duopoly model where sellers choose algorithms. An algorithm is a mapping from the competitor's price to own price. Once set, algorithms respond quickly. Customers arrive randomly and so do opportunities to revise the algorithm. In the simple game with two possible prices, monopoly pricing is the unique equilibrium outcome for standard functional forms of the profit function. More generally, with multiple prices, market power is the rule–in all equilibria, the expected payoff of both sellers is above the competitive outcome, and that of at least one seller is eventually close to or above the monopoly outcome. Sustenance of such collusion seems outside the scope of standard antitrust laws for it does not involve any direct communication.

1 Introduction

There is growing concern amongst academics, policymakers, and increasingly consumers that algorithms are engendering otherwise outlawed collusive practices amongst firms in various industries. For example, writing in *Science* magazine, Calvano, Calzolari, Denicolò, Harrington Jr., and Pastorello [2020a] remark:

> Collusion is generally condemned by economists and policy-makers and is unlawful in almost all countries. But the increasing delegation of price-setting to algorithms (1) has the potential for opening a back door through which firms could collude lawfully (2). Such algorithmic collusion can occur when artificial intelligence (AI) algorithms learn to adopt collusive pricing rules without human intervention, oversight, or even knowledge.

There is growing empirical evidence for this too. Studying adoption of algorithmic pricing in German gas stations, Assada, Clark, Ershovc, and Xud [2021] find

> Adoption increases margins, but only for non-monopoly stations. In duopoly markets, margins increase only if both stations adopt, suggesting that AP has a significant effect on competition.

Finally, new experimental work also suggests that algorithms may lead to collusion using adaptive strategies oft invoked in dynamic games. Calvano, Calzolari, Denicolò, and Pastorello [2020b] write:

> The high prices are sustained by collusive strategies with a finite phase of punishment followed by a gradual return to cooperation. This finding is robust to asymmetries in cost or demand, changes in the number of players, and various forms of uncertainty.

The objective of this paper is to offer a parsimonious model that captures the essence of why algorithms may be abetting collusion. We focus on the two specific properties of algorithms (also outlined in Brown and MacKay [2021]). First, by design they allow a quick response time to price changes by competitors. Second, they offer temporary commitment to a pricing strategy—once an algorithm is chosen, price setting is delegated to it, till the algorithm is updated. The paper argues that these properties indeed facilitate collusion— the market *naturally* converges to prices higher than the competitive outcome. Moreover, these choices looks legal on the books since they does not require any direct communication.

More specifically, we study a duopoly model of price setting, where the strategy for each seller is to choose an algorithm. An algorithm is simply a mapping from competitor's price to own price. Once the two algorithms are chosen, they automatically set prices in response to each other. Customers arrive according to some Poisson process. Upon arrival, a probabilistic demand function determines which seller the customer may buy from—a lower price leads to a higher likelihood of being chosen. In addition, each seller gets an opportunity to revise their algorithm according to a separate Poisson process. The dynamic game essentially follows a

sequence of customer arrivals and opportunities to revise algorithms (see Section 2 for the model).

The key modeling assumption is that "Δt", i.e. the unit that measures the reaction time of algorithms, is small. It delivers three mechanisms: First, sellers can experiment and learn the opponent's algorithm very quickly, the learning cost is at most linear in the set of prices and the unit of time. So effectively when they choose their algorithm, they are best responding to the algorithm of the opponent. Second, prices converge quickly for a fixed set of algorithms and the convergent prices are what essentially matter for accounting payoffs. Third, since the time period of price changes is smaller than the average time sellers have to change algorithms, this offers them temporary commitment to a pricing strategy, in contrast to a standard repeated game. Once two algorithms are set, they react mechanically to the opponent's price, till one of the sellers gets a new opportunity revise its algorithm.

In fact, the first step in the analysis is to show that this dynamic game reduces to an asynchronous repeated game, that is, a repeated game in which players rotate in picking strategies. We then focus on Markov perfect equilibria of this game. The main result of the paper is that some level of collusion is inevitable, i.e. it is always attained, for all specifications of the model, and that these collusive equilibria can be further characterized. For starters, the payoff of both sellers is greater than the competitive outcome.[1] Moreover, exercise of market power is omnipresent—in all equilibria, the payoff of at least one seller is bounded from below by a constant that approximates the monopoly profit.[2] Further still, for certain specifications, repeated play of the monopoly strategy is the unique convergent outcome. In short, collusion is inevitable and extreme collusion is commonplace.

The model is closely related to Salcedo [2015], which presents a remarkably elegant way to capture the idea of collusion through algorithms. Salcedo [2015], however, characterizes subgame perfect equilibria and asserts the main result for the limit case of zero effective discounting, that is zero revision opportunities for the sellers, which we argue is unrealistic. Our tack on the problem delivers results for all values of effective discounting, and in addition presents a much more exhaustive characterization of (Markov) equilibrium outcomes. A detailed discussion is presented in Section 6.

Our analysis starts with a special case where the sellers have access to two prices— the competitive one and the pure monopoly one (Section 3). Within that realm, we consider a general 2×2 matrix game that satisfies standard restrictions for the Prisoner's Dilemma game. It is shown that for two standard specifications–when the profit function is either quadratic or takes the discrete choice form used in various applied models– repeated play of the monopoly price is the unique equilibrium outcome. Further, for general profit functions too, three types of equilibria emerge, two of which lead to monopoly as the unique outcome and the payoffs in the third too are bounded away from the competitive outcome. Recollect that repetition of static Nash prices is always an equilibrium in the standard repeated games model, not here.

---

[1]By "competitive outcome", we mean here the Nash equilibrium of the one-shot game.

[2]In a duopoly model, we abuse terminology slightly and term monopoly profit to be half of the profit that would accrue to a single seller who sets prices to maximize its profit.

Here is a brief intuition for the inevitability of collusion in the 2×2 model: The universe of all algorithms can be reduced to four types. A seller could always chooses the competitive price, no matter what the opponent does, call this algorithm/strategy s_C; the seller can also always choose the monopoly price, call this s_M. In addition, the seller can play tit-for-tat, that is, choose whatever price the opponent chooses, call this s_T or it can play reverse-tit-for-tat, which means picking the opposite of what the opponent chooses, call this s_R.

Suppose seller A chooses s_C, that is, always pick the competitive price. Then, upon their turn to set an algorithm, seller B can also pick s_C so that the competitive price is selected by both. In a standard repeated game this would result in repetition of static Nash, which is of course an equilibrium outcome. But in this set up, is it optimal for B to choose s_C? It is not: if instead B chose s_T, then the competitive outcome will be played till A has an opportunity to update their algorithm from s_C. However, once A has a chance to revise their algorithm, the best response is also to choose s_T and start off the new phase with the monopoly price. Since both are playing s_T, the monopoly price is selected repeatedly, which generates a higher payoff. This points towards why repetition of the competitive price cannot be an equilibrium outcome. The repeated choice of s_C does not satisfy the perfectness criterion in Markov perfect equilibrium.

Building on this logic further by comparing 4x4 matches of all algorithms, one can in fact prove that the repeatedly play of the monopoly price emerges globally (i.e. for all parameters) as an equilibrium outcome. Moreover, for many reasonable specifications of payoffs, it is the unique outcome, hence the claim of extreme collusion being commonplace.

We then take these insights to a general model with any finite grid of prices (Section 4). This of course complicates the analysis (at least theoretically) by orders of magnitude. Even with three possible prices there are 27 algorithms to choose from and with four there are 256 and so on. Despite this complexity we draw a general theoretical lesson. It says that payoffs in all Markov perfect equilibrium are such that: (i) both sellers receive a payoff higher than the competitive outcome, and (ii) at least one sellers' payoff is bounded from below by a constant that approximates monopoly profit.

At a high level, the intuition for this result is similar to the 2x2 model. The low price equilibria are "learnt away" by following a chain of temporary commitments to a higher price. Within each commitment phase of set algorithms, prices converges quickly because of the smallness of response times, so only the convergent outcomes are taken to be payoff relevant. And while setting algorithms, sellers take into account future opportunities for the competitor to revise their algorithms and so in the larger repeated game, algorithms too eventually settle on to a stable pair or cycle of pairs.

In the 2x2 model, a two step chain took us from one seller choosing s_C to both eventually choosing s_T, which sustained the monopoly outcome. With multiple possible prices, more complex chains can be formed. We consider a subset of these to show that in any equilibrium at least one seller will end up guaranteeing themselves a payoff which is bounded from below by a high enough constant.

To prove the main results, we make simplifying assumptions on how payoffs are calculated and how cycles (as opposed to convergent prices) are evaluated for a fixed set of algorithms. These are relaxed in extensions (Section 5) to show that the basic insights carry through. We also extend equilibrium notion to sugbame perfect and show that a much larger class of equilibria emerge where both collusion and conflict can be sustained. Moreover, the specific sequence of algorithms required to sustain these payoffs are quite complex.

In terms of techniques, the paper draws on the literature on asynchronous repeated games, especially Maskin and Tirole [1988a,b]. These set of papers explore temporary commitment by fixing a player's strategy for some finite amount of time and rotating the opportunity to revise them, hence the term 'asynchronous' (see also Lagunoff and Matsui [1997]). The technical analysis is also related to repeated games with renegotiation (see for example Pearce [1987], Miller and Watson [2013], Safronov and Strulovici [2018]) in that opportunities to set algorithms can be seen as renegotiation constraints. In addition, algorithms carry the flavor of finite automaton in repeated games, as for example in Abreu and Rubinstein [1988]. Most of this literature however has players setting their automaton at the beginning with no recourse for revisions. Finally, in allowing for random opportunities to revise strategies, we also relate to the literature on revision games, see for example Kamada and Kandori [2020].

We of course speak to a now burgeoning literature on the foundations and implications of algorithmic pricing. Brown and MacKay [2021] document and then model asymmetries in price dispersion arising out of algorithmic pricing. Asker, Fershtman, and Pakes [2021] study the impact of different learning protocols used by artificial intelligence on algorithmic pricing. Johnson, Rhodes, and Wildenbeest [2021] explore experimentally the ability of a platform to increase consumer welfare by reducing prices in differentiated product models where sellers use learning algorithms. Harrington Jr. [2022] presents a model in which sellers outsource pricing to a third party that uses an algorithm and finds that the algorithm does not reduce competition but it makes prices more sensitive to underlying variation in demand. Hansen, Misra, and Pai [2021] argue that when competing sellers are using algorithms to learn some features of demand, collusion emerges because the sellers end up running correlated experiments.

The study of collusion (or more broadly cooperation) has been a central question in much of game theory (as far back as Aumann [1959] and Friedman [1971]). Green and Porter [1984] argued that simply observing the competitors' prices, as opposed to entire strategies, is sufficient to sustain collusion. More recently, Awaya and Krishna [2016] show that communication can in fact help sustain collusion even if firms may not be able to track others' prices. Sugaya [2020] proves a meta folk theorem for repeated games with private monitoring. In comparison to these, we keep the structure of information simple in that all actions are observable, but emphasize that quick reaction and temporary commitment offered by algorithms makes collusion not just possible; in fact some level of collusion is inevitable and extreme collusion is commonplace.[3]

---

[3]There is also a significant applied and empirical literature on the economics of detecting and understanding mechanics of collusion, especially in auctions. See Asker [2010] and Chassang, Kawai, Nakabayashi, and Ortner

While the model is stylized, it captures some essence of why automated price settings can lead to sustained markups that regulators can no longer ignore. Indeed, the endeavor here is to tease out sustained monopoly as the unique (in the two-price model) outcome or at least extreme collusion as the frequent (in the multi-price model) outcome, so that the theoretical point can be made starkly.

The challenge for anti-trust authorities assumes great significance for old methods of defining, detecting and hence describing collusion in a court of law are limitedly effective. If collusion can be sustained without explicit communication, how do we proceed? In fact, Harrington Jr. [2018] argues that "collusion by software programs which choose pricing rules without any human intervention is not a violation" of standard anti-trust laws used in the United States. This is a puzzle that several antitrust and competition commissions across the world are now encountering or will do in the near future, which is an invitation for many more theoretical and empirical studies to inform law and policy.

2 Model

2.1 Market and prices

Two sellers $i = A, B$ are selling a product. Each sets a price p^A and p^B. The marginal costs are normalized to 0. Customers arrive at a Poisson rate λ, so the probability of having a customer in the interval Δt is approximately $\lambda \cdot \Delta t$. Each customer purchases at most one unit of the product.

Demand is defined probabilistically: At current prices p and q, the probability that a customer purchases the product from the firm with price p is given by $d(p, q)$. So, total demand is given by $d(p, q) + d(q, p)$. Fix the payoff function to be

$$\pi(p, q) = p \cdot d(p, q). \tag{1}$$

Then, the expected payoff from one customer for each seller is given by

$$\begin{aligned} \pi^A &= \pi(p^A, p^B) \\ \pi^B &= \pi(p^B, p^A) \end{aligned} \tag{2}$$

We will impose the following assumptions om the payoff function $\pi(p, q)$:

Assumption 0. (i) $\pi(p, q)$ is increasing in q and for each q there is a unique global maximum of $\pi(p, q)$ in p which is increasing q; (ii) there is a unique static Nash equilibrium with prices (p_C, p_C) (where C stands for "competitive" price); (iii) the joint profit $\pi(p, q) + \pi(q, p)$ is maximized at $p = q = p_M > p_C$ (where M stands for "monopoly" price); (iv) for any $p \geqslant \arg\max_{p'} \pi(p', p_M)$, the point $(\pi(p, p_M), \pi(p_M, p))$ is on the boundary of the convex hull of

---

[2022] for some recent work, and Marshall and Marx [2012] for a textbook treatment. Collusion due to the use of algorithms though is a relatively new addition to this cannon. In recent work, Banchio and Skrzypacz [2022] explore implications of simple learning algorithms for the sustenance of collusion in auctions.

6

the payoffs set $\{(\pi(p,q), \pi(q,p)) \ for \ p, q \in [p_C, p_M]\}$.

Parts (i) to (iii) are standard assumptions that correspondent π to a standard Bertrand duopoly. Part (iv) is a useful technical assumption that helps us narrow down the number of deviations to consider when stating the general result (Theorem 2). Several commonly used market demand functions (including the linear demand and discrete choice demand mentioned in Corollary 2) satisfy these assumptions.

We also restrict the set of possible prices for each seller to a discrete subset of equally spaced numbers between the competitive p_C and monopoly prices p_M:

$$p_C = p_1 < \cdots < p_K = p_M. \tag{3}$$

Fix $\Delta = p_i - p_{i-1}$ for $i = 1, \ldots, K$ to be the difference between two consecutive prices. In a slight abuse of notation, we denote:

$$
\begin{aligned}
p_{M-\Delta} &= p_M - \Delta = p_{K-1} \\
p_{M-2\Delta} &= p_M - 2 \cdot \Delta = p_{K-2} \\
&\cdots \\
p_{C+\Delta} &= p_C + \Delta = p_1
\end{aligned}
\tag{4}
$$

2.2 Algorithms

An algorithm, in its simplest form, is a function $s : p^{-i} \to p^i$ that determines which price p^i is chosen by seller i in response to the competitor's price p^{-i}.[4] Prices are observable but sales of the competitor are not, or equivalently algorithms react only to prevailing prices, which we think is a reasonable assumption, especially for online sellers.

We assume algorithms can change prices once per small interval of time Δt. Suppose, the prices react at the following calendar times: $0, \Delta t, \ 2 \cdot \Delta t, \ 3 \cdot \Delta t, \ldots$ If s^A and s^B are the corresponding algorithms of the two sellers, then the prices will evolve according to

$$
\begin{aligned}
p^A_{t+\Delta t} &= s^A \left(p^B_t \right) \\
p^B_{t+\Delta t} &= s^B \left(p^A_t \right)
\end{aligned}
\tag{5}
$$

If the algorithms of the two sellers are temporarily fixed, then there are only two possible outcome for the price sequences generated from (5).

1. They converge to fixed price pair (p^A, p^B) such that:

$$p^A = s^A(p^B) \qquad p^B = s^B(p^A) \tag{6}$$

---

[4] Algorithms broadly are a set of instructions that mechanically map some input information to an pre-specified set of outputs. Different notions of algorithms have been adopted thus far in the economics literature. The definition used here mostly follows Brown and MacKay [2021]. In Asker et al. [2021] and Johnson et al. [2021], sellers adopt the Q-learning algorithm, often used in practice. In Harrington Jr. [2022], an algorithm maps demand shocks not observable to the sellers to a prescribed price. A generalized version of our notion, also used in Salcedo [2015], defines an algorithm to be a mapping $s : \omega^i, p^{-i} \to \nu^i, p^i$, where ω^i represents the current state of the algorithm of seller i and ν^i is the subsequent state. As discussed later in Section 7, if the cardinatlity of these states is not too large, our results can largely be extended to such algorithms.

2. They converge to a cycle of length T:

$$
\begin{aligned}
p_1^A \to\ & p_2^A \to\ \cdots \to\ p_T^A \to\ p_1^A \to \ldots \\
p_1^B \to\ & p_2^B \to\ \cdots \to\ p_T^B \to\ p_1^B \to \ldots
\end{aligned}
\tag{7}
$$

The sellers can also periodically adjust their algorithms. The possibility to adjust the algorithms arrives at an exogenous Poisson process with intensity μ; so the probability of the event happening in the interval Δt is approximately $\mu \cdot \Delta t$. In addition, whenever one of the sellers adjusts their algorithm, they can pick an initial short sequences of prices. So, for a few subsequent intervals of length Δt, prices can be chosen by the seller intentionally, after which the price setting is delegated to the new algorithm.

This initial set of prices can be employed for two objectives. First, sellers experiment and learn the opponent's algorithm before setting their own. It takes at most K prices adjustments (K is the cardinality of price set) to fully understand the opponent's algorithm. The cost of this is proportional to $K \cdot \Delta t$, which is linear in K and unit of time, and hence small for a small Δt. Second, upon deciphering the opponent's algorithm, the seller adjusting the algorithm can choose initial few prices to select their most preferred price sequence that satisfies (5) for the prevailing algorithms s^A and s^B. Think of this as the initial condition for the new system that is then delegated to determine prices. Assumption 1 below ensures that the initial prices for both these objectives, while important in setting algorithms, are not payoff relevant.[5]

2.3 Simplifying assumptions

In what follows, we make three simplifying assumptions. These make the overall analysis tractable and allow us to focus on the "big picture" of how quickness of response and temporary commitment offered by algorithms leads to collusion. In section 5 we relax two of these and show that our key results continue to hold.

- **Assumption 1.** The interval Δt is sufficiently small (in particular, we need $r\Delta t \ll 1$ and $\mu\Delta t \ll 1$) so that initial price adjustments governed by (5) can be ignored and only convergent prices, given by (6) or (7) are payoff relevant.

- **Assumption 2.** Cycles are ruled out. A seller changing their algorithms is not allowed to choose a new algorithm (and an initial price sequence) such that together with the existing algorithm of the opponent the pair results in a cycle (7).

- **Assumption 3.** If a seller adjusting their algorithm is indifferent between two options (in terms of continuation payoffs), the seller would choose the one that is preferred by their opponent.

---

[5]A different interpretation of similar modeling assumptions is used in Salcedo [2015] (and to some extent in Brown and MacKay [2021] as well). There the time between revision opportunities is used by the seller to decipher the opponent's algorithm. So, learning is captured exogenously through the random arrival times to set algorithms. The same interpretation can be applied in our setting. The experimentation interpretation, which we slightly prefer, is facilitated by the smallness of time required for price adjustments in our model. In Appendix E of the paper, we discuss the issue of experimentation further.

Assumption 1 represents the two key features of algorithms that we attempt to model, that they respond quickly, and they offer temporary commitment to sellers. These are ensured by $r\Delta t \ll 1$ and $\frac{1}{\mu} \gg \Delta t$. As a consequence, only the convergent prices matter for payoffs. We view the smallness of Δt as an important aspect of algorithmic pricing and thus a critical modeling choice for our analysis. The assumption that adjustments driven by (5) pre-convergence can be ignored is made for simplicity and will be relaxed in Section 5 to show that all the basic insights carry through.

Our motivation for **assumption 2** is that you normally would not see cycles in reality. A cycle with frequently changing prices (for example, every minute or even every hour) would be confusing for consumers. Moreover, when consumers see frequently changing prices they can time their purchases and wait for the best prices from each of the sellers, which is not in the sellers' interest either.[6] In Section 5 we show that our results for the case of two possible prices do not change if we do indeed allow cycles.

Assumption 3 is typically never binding. If two different algorithms generate different subsequent paths in terms of prices, then the continuation payoffs would only be identical for certain very special discount factors β or for certain very special payoff matrices $\{\pi(p,q)|p,q = p_C, p_{C+\Delta}, \dots, p_M\}$. We make this assumption so that even in such special cases we do not have to deals with the multiplicity of optimal responses.

2.4 Expected payoffs and reduction to an asynchronous repeated game

Suppose that the current algorithm of seller i is s^i and the current algorithm of their competitor is s^{-i} and seller i was the last to adjust their algorithm. Denote by

- $\xi(s^{-i}, s^i) = (p^{-i}, p^i)$: the price pair to which algorithms would converge; it would be a price pair consistent with these two algorithms and since i was the last to adjust their algorithm, amongst all such price pairs it will be the one most preferred by i.

- $\pi(s^i, s^{-i}) = \pi(p^i, p^{-i})$: the expected payoff from one customer for the seller who was the last to adjust their algorithm after we reach the convergent price pair.

- $\bar{\pi}(s^{-i}, s^i) = \pi(p_{-i}, p_i)$: the expected payoff from one customer for the seller who was *not* the last to adjust their algorithm after we reach the convergent price pair.

Index all the moments when one of the sellers is adjusting their algorithm with $k = 0, 1, 2, \dots$. Suppose the seller adjusting their algorithm at moment k chooses the algorithm s_k, thus we end up with the following sequence of algorithms $s_0, s_1, s_2, \dots$. In this sequence all "even" algorithms correspond to one seller and all "odd" algorithms correspond to the other.[7]

---

[6]Formally, disallowing cycles amounts to restricting the strategy space of the underlying stage game, that is, for a fixed strategy of the opponent, a restricted set of strategies is available to a seller. Since the probability of both sellers receiving opportunities to revise the algorithm at the same time is zero, **Assumption 2** still leads to a well defined dynamic game. One can rationalize this through a high cost of entering a cycle.

[7]Suppose seller A gets the first random draw to set their algorithm. Then even if they get a second random draw to adjust their algorithm before seller B's first turn, it is never optimal to deviate from the current algorithm.

Now consider some arbitrary adjustment moment k. Denote by U_k the expected continuation payoff of the seller adjusting their algorithm, and by V_k the expected continuation payoff of the competitor. With Assumption 1, we can calculate the continuation payoffs recursively:

$$U_k = \int_0^\infty e^{-rt} \cdot e^{-\mu t} \left[\lambda \cdot dt \cdot \pi(s_k, s_{k-1}) + \mu \cdot dt \cdot V_{k+1} \right] = \frac{\lambda}{r+\mu} \pi(s_k, s_{k-1}) + \frac{\mu}{r+\mu} V_{k+1} \tag{8}$$

$$V_k = \int_0^\infty e^{-rt} \cdot e^{-\mu t} \left[\lambda \cdot dt \cdot \bar{\pi}(s_{k-1}, s_k) + \mu \cdot dt \cdot U_{k+1} \right] = \frac{\lambda}{r+\mu} \bar{\pi}(s_{k-1}, s_k) + \frac{\mu}{r+\mu} U_{k+1} \tag{9}$$

If we normalize the continuation payoffs as

$$u_k = \frac{U_k}{\lambda/r} \text{ and } v_k = \frac{V_k}{\lambda/r}, \tag{10}$$

and denote normalized discount factor as

$$\beta = \frac{\mu}{r+\mu} \tag{11}$$

then we get the following two simple equations describing the evolution of payoffs:

$$\begin{aligned} u_k &= (1-\beta) \cdot \pi(s_k, s_{k-1}) + \beta \cdot v_{k+1} \\ v_k &= (1-\beta) \cdot \bar{\pi}(s_{k-1}, s_k) + \beta \cdot u_{k+1} \end{aligned} \tag{12}$$

Thus, we have essentially converted our more complicated 'continuous time game' into an *asynchronous repeated game* in discrete time with discount factor β, in which one seller adjusts their action (which is an algorithm) in "even" periods and the other in "odd" periods and the payoffs in each in period are determined by finite matrices $\pi(s_k, s_{k-1})$ and $\bar{\pi}(s_{k-1}, s_k)$.

A comment on the normalization above is in order. Note that the normalized expected payoffs from equations (12) do not directly depend on customer arrival rate λ. This happens because the rate is essentially a scaling factor that multiplies profits of each seller, and, thus does not affect their optimal behavior. The effective discount fact β is determined by the ratio between the frequency of algorithm revision opportunities μ and the interest rate r faced by the sellers. In realistic scenarios $r \ll \mu$, and, thus, β is typically close to 1, but still bounced away from it. For example, if the interest rate is 5% per year and each seller has (on average) 5 opportunities to revise their algorithm in a year, then $\beta = \frac{5}{0.05+5} \approx 0.99$.

2.5 Markov perfect equilibria

In our analysis we focus mainly on Markov Perfect Equilibria. In such equilibria when one of the sellers is adjusting their algorithm, their best response depends only on the current algorithms used by their opponent. This means that if the seller would face the same algorithms of the opponent in the future, they would respond in the same manner. In Section 5 we discuss the more general class of subgame perfect equilibria.

Two quick things to note in motivating MPE: First, while equilibria with richer intermediate levels of collusion are possible with subgame as opposed to Markov perfection, they are typically rather complicated. They require each seller to have very specific beliefs about the

behavior of their competitor, which may not hold muster in reality. Second, we want to point out that since the space of possible algorithms is typically quite large, even Markov Perfect Equilibria are not necessarily very simple. The reader will shortly learn these equilibira still cover a wide range of possibilities.[8]

In a Markov Perfect Equilibrium, the current algorithm of the competitor, say s, is regarded as the state variable. Then, for each seller ($i = A, B$) we can define:

- $f^i(s)$: best response of seller i to the strategy s of the opponent when i is adjusting their algorithm.

- $u^i(s)$: the continuation payoff of i if i is changing their algorithms and the current algorithms of the opponent is s.

- $v^{-i}(s)$: the continuation payoff of $-i$ if i is changing their algorithm and the current algorithm of $-i$ is s.

In equilibrium the best response $f^i(s)$ and the continuation payoffs u^i and v^{-i} should satisfy:

$$u^i(s) = \max_{s'}(1 - \beta) \cdot \pi(s', s) + \beta \cdot v^i(s')$$

$$f^i(s) \in \arg\max_{s'}(1 - \beta) \cdot \pi(s', s) + \beta \cdot v^i(s') \tag{13}$$

$$v^{-i}(s) = (1 - \beta) \cdot \bar{\pi}(s, f^i(s)) + \beta \cdot u^{-i}(f^i(s))$$

The sequence of algorithms $s_0, s_1, s_2, \ldots$ in which each algorithm is a best response to the previous one (i.e. for any k, $s_{2k+1} = f^i(s_{2k})$ and $s_{2k+2} = f^{-i}(s_{2k+1})$) will be termed an *equilibrium sequence* starting with s_0.

3 Two prices and reduction to a Prisoner's dilemma game

3.1 Expected payoff from a single customer

We start our analysis with a simple case in which each seller $i = A, B$ can choose one of the two prices: p_M (monopoly) or p_C (competitive). The expected payoffs from one customer can be described with a simple two-by-two matrix (see Table 1). Assumption 0 then reduces to the following conditions.

The payoff matrix has the Prisoners' dilemma type structure, that is

$$\pi(p_M, p_C) < \pi(p_C, p_C) < \pi(p_M, p_M) < \pi(p_C, p_M), \tag{14}$$

---

[8]Maskin and Tirole [1988b] (in Section 10 of the paper) provide four reasons for studying MPE as a benchmark of sorts in repeated oligopoly models, and we refer the reader to their persuasive pitch for our choice of MPE. To be clear, this is not to say that characterizing SPE is not important, just that MPE is a good benchmark to work with. Moreover, it is perhaps more surprising that extreme collusion is widespread in MPE, that varying levels of collusion can be sustained in SPE for high discounting is to be expected.

$A \backslash B$	p_M	p_C
p_M	$\pi(p_M, p_M)$, $\pi(p_M, p_M)$	$\pi(p_M, p_C)$, $\pi(p_C, p_M)$
p_C	$\pi(p_C, p_M)$, $\pi(p_M, p_C)$	$\pi(p_C, p_C)$, $\pi(p_C, p_C)$

Table 1: Expected payoffs from one customer for the two prices model

$A \backslash B$	p_M	p_C
p_M	1, 1	$-y$, $1 + x$
p_C	$1 + x$, $-y$	0, 0

Table 2: Normalizing and re-writing the payoff matrix in Table 1 in terms of x and y.

In addition, the monopoly outcome (p_M, p_M) is superior from a utilitarian perspective

$$2 \cdot \pi(p_M, p_M) > \pi(p_C, p_M) + \pi(p_M, p_C). \tag{15}$$

In order to describe how possible equilibria depend on the payoff matrix, it would be convenient to define

$$x = \frac{\pi(p_C, p_M) - \pi(p_M, p_M)}{\pi(p_M, p_M) - \pi(p_C, p_C)} \qquad y = \frac{\pi(p_C, p_C) - \pi(p_M, p_C)}{\pi(p_M, p_M) - \pi(p_C, p_C)}. \tag{16}$$

Then, the Prisoners' dilemma game in Table 2 is "equivalent" to the original game in the sense that set of equilibria in both Tables 1 and 2 are exactly the same. So, we can describe all possible equilibria in the static game in terms of x and y, and in the dynamic game in terms of x, y and β.

The parameter x shows by how much the deviation from monopoly, $\pi(p_C, p_M)$, exceeds the monopoly payoff, $\pi(p_M, p_M)$, and y shows by how much $\pi(p_M, p_C)$ is below the competitive payoff, $\pi(p_C, p_C)$. Both are measured relative to the total distance $\pi(p_M, p_M) - \pi(p_C, p_C)$. Conditions (14) and (15) are now equivalent to

$$\begin{aligned} x, y &> 0 \\ y &> x - 1 \end{aligned} \tag{17}$$

3.2 Algorithms and equilibria

With two prices we can easily describe all possible algorithms that the sellers could use, see Table 3. Here, s_C represents the algorithm that chooses price p_C no matter what the competitor does, and similarly s_M is the algorithm that always chooses p_M. Further, s_T is the tit-for-tat algorithm that replicates the competitor's action, and s_R is the reverse-tit-for-tat that chooses the exact opposite of the competitor's action.

Given four possible choices of algorithms, we can completely describe convergent outcomes, this is reported in Table 4. The entries in the table report the possible final prices for any pairs of algorithms, $\xi(s^{-i}, s^i)$. For example, when one of the algorithms is s_T and the other is s_R we end up with a cycle (which is ruled out as a final outcome by assumption). When both

s_C	$p_C, p_M \to p_C$	always "competitive" price
s_M	$p_C, p_M \to p_M$	always "monopoly" price
s_T	$p_C \to p_C, p_M \to p_M$	tit-for-tat
s_R	$p_C \to p_M, p_M \to p_C$	"reverse" tit-for-tat

Table 3: Possible algorithms for two prices

	s_M	s_C	s_T	s_R
s_M	p_M, p_M	p_M, p_C	p_M, p_M	p_M, p_C
s_C	p_C, p_M	p_C, p_C	p_C, p_C	p_C, p_M
s_T	p_M, p_M	p_C, p_C	p_M, p_M	$p_C, p_C \to p_C, p_M \to p_M, p_M \to p_M, p_C$
s_R	p_C, p_M	p_M, p_C	$p_C, p_C \to p_M, p_C \to p_M, p_M \to p_C, p_M$	p_C, p_M

Table 4: Final outcomes for different pairs of algorithms assuming row player was the last to choose algorithm

algorithms are s_R, the seller who was the last to adjust their algorithm can the choose initial price sequence and, thus, out of the two possible final prices pairs (p_C, p_M) and (p_M, p_C) we choose the one that is preferred by this seller.

Now, we can state the main result that completely characterizes equilibria for this two-price game. For starters, an equilibrium always exists. In addition, the universe of equilibria takes a simple form in that it all boils down to three parameters: x, y and β, and we can characterize this set precisely. To simplify notation, we write $s^{-i} \to s_i$ to mean $f(s^{-i}) = s^i$, that is the strategy s_i optimally chosen by seller i when the competitor is playing s^{-i}.

Theorem 1.

1. *A Markov equilibrium always exists.*

2. *If $x \leq \beta$ only the following Markov equilibrium is possible:*

$$
\begin{aligned}
s_R &\to s_C \to s_T \\
s_T &\to \{s_T, s_M\} \qquad (type\ I) \\
s_M &\to \{s_T, s_M\}
\end{aligned}
$$

3. *If $x > \beta$ then*

   (a) *the following equilibrium is always possible:*

$$
s_R, s_M \to s_C \to s_T \to s_T \qquad (type\ II)
$$

   (b) *in addition if $y < \beta \cdot (x - \beta)$ the following equilibrium is also possible:*

$$
\begin{aligned}
s_M, s_R &\to s_R \\
s_C, s_T &\to s_T
\end{aligned} \qquad (type\ III)
$$

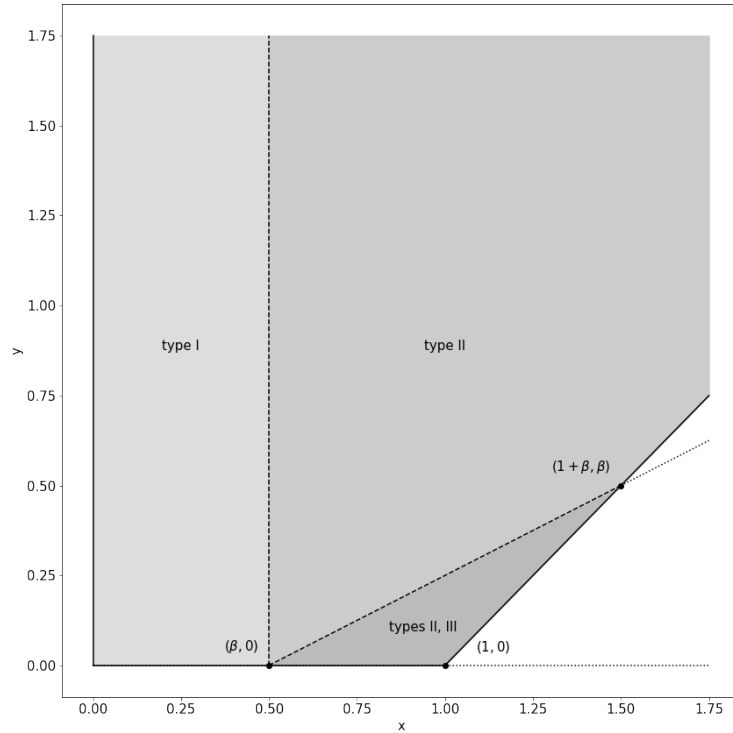Theorem 1 describes which equilibria are possible depending on structure of the payoff

13

Figure 1: Possible equilibria for varying values of x and y, β fixed at 0.5.

matrix (i.e., x and y) and effective discount factor (i.e. β). For example, if β is lower than x, then in the type II equilibrium, for algorithm s_R or s_M chosen by the opponent, it is optimal for the seller to set the algorithm s_C. This is because for a low enough discounting, the player choosing the algorithm cares about the immediate payoff and thus wants prices to converge to (p_C, p_M) in the short run. However, going further, when it is the opponent's turn to react to s_C, they choose s_T for reasons explained in the introduction; and so on. Figure 1 illustrates the result graphically—fixing β it partitions the (x, y) space into the three types of equilibria reported in Theorem 1.

The proof of Theorem 1 takes a general approach as it treads through the set of best responses for each chosen algorithm of the opponent. We first show that it is always optimal to respond to s_C with s_T. To do that we consider alternative responses and equilibrium sequences of algorithms that such responses can generate and show that in each of them it would be optimal for at least one of the sellers to deviate from the current strategy. Then, we show after s_T, the best response is either s_T or s_M and we can only observe monopoly prices (p_M, p_M) even after subsequent opportunities to revise algorithms. Finally, we find the best responses to s_R and s_M by comparing the expected payoffs for each of the alternatives.

What are the pricing outcomes corresponding to these equilibria? The next result documents that the answer to this follows immediately from matching the three types of equilibria reported in Proposition 1 with possible final outcomes reported in Table 4. It states that for equilibria for types I and II, the unique outcome is the monopoly price (p_M, p_M), and for type III equilibrium prices oscillate between (p_C, p_M) and (p_M, p_C).
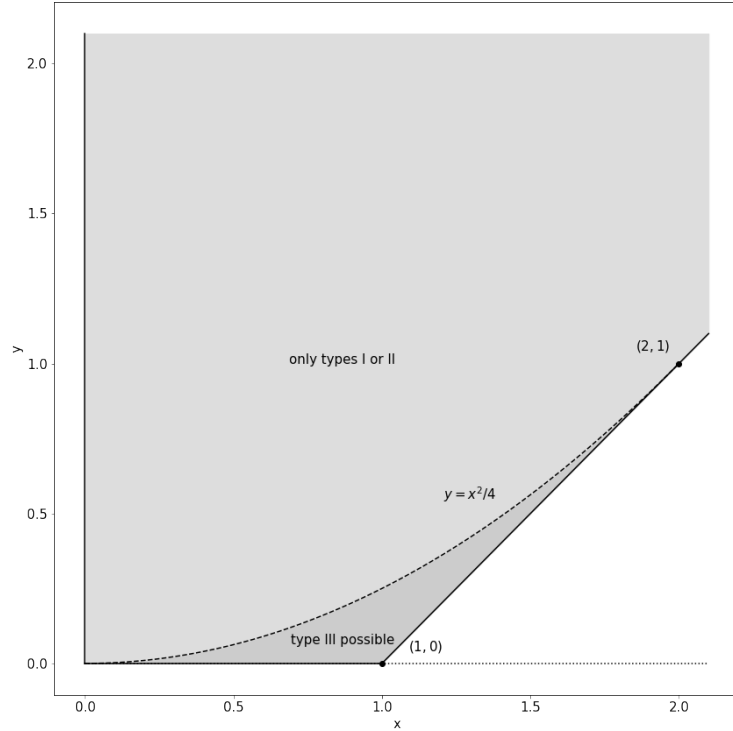
14

Figure 2: When type III equilibrium with alternating prices can arise for some value of β.

Corollary 1.

1. *In any equilibrium we either converge to a monopoly price pair (p_M, p_M) (for equilibria of types I or II) or end up with an alternating sequence of prices*

$$(p_C, p_M) \to (p_M, p_C) \to (p_C, p_M) \to \dots \tag{18}$$

in which prices change every time one of the sellers gets the opportunity to adjust their algorithm (for equilibrium of type III).

2. *Type III equilibrium (and the corresponding alternating price sequence) is only possible if*

$$x < 2, \qquad y < \frac{x^2}{4}, \ and \qquad \frac{x}{2} - \sqrt{\frac{x^2}{4} - y} < \beta < \frac{x}{2} + \sqrt{\frac{x^2}{4} - y}.$$

The main message in Corollary 1 is that some level of market power is inevitable, for (p_C, p_C) is never a convergent equilibrium outcome. Moreover, in *most* cases, the repeatedly play of the pure monopoly price is the unique equilibrium outcome. This claim can be visualized in Figure 2, which shades the values of x and y for which type III equilibrium is possible at least for some values of the discount factor β. Note that "for some β" is a fairly weak criterion, and still the region is quite small. This leads us to conclude that equilibria of types I and II dominate the parameter space and hence the unique monopoly outcome is to expected in such two-price games.

As a final step, we provide some simple sufficient conditions on the payoff matrix and the

15

profit function under which repeated play of the monopoly price is the unique equilibrium outcome.

Corollary 2.

1. *If*

$$y > x \quad \text{that is} \quad \pi(p_M, p_C) + \pi(p_C, p_M) < \pi(p_M, p_M) + \pi(p_C, p_C)$$

then only equilibria of type I or II are possible, which means that prices always converge to a monopoly outcome (p_M, p_M).

2. *If payoff function $\pi(p, q)$ is quadratic (demand is linear):*

$$\pi(p, q) = p \cdot \frac{D - p + \alpha(q - p)}{2D}$$

or if π is a "discrete choice" profit function:

$$\pi(p, q) = p \cdot \frac{e^{-bp}}{a + e^{-bp} + e^{-bq}}$$

then again in any equilibrium, prices always converge to a monopoly outcome (p_M, p_M).

That $y > x$ leads uniquely to the monopoly outcome is already visible in Figure 2. Conceptually, a quick look at Table 2 tells us that if the gain from "deviating" from (p_M, p_M) is smaller than the loss the opponent would incur from this deviation, then the opponent cannot be incentivized with continuation value for any value of β.

Part 2 of the result shows that for two widely used profit functions—the linear demand case invoked in various textbooks and the discrete choice model used often in applied work, repeatedly play of the monopoly outcome is again the unique equilibrium outcome. The proof essentially shows that for these profit functions the condition $y > x$ is always satisfied, hence only type I or II equilibria arise.

4 Multiple prices

4.1 Main result

In this section we consider the case in which for each seller there are more than two possible prices: $p_C, p_{C+\Delta}, \ldots, p_{M-\Delta}, p_M$. With multiple prices the number of possible algorithms explodes (there are $3^3 = 27$ algorithms for 3 prices, and $4^4 = 256$ algorithms for 4 prices and so on). Thus, it is now difficult to "brute force" our way through all best response functions to describe Markov equilibria.

With multiple prices there also exist *unequal* equilibria, in which one seller makes a profit even higher than the monopoly outcome. Suppose (p, q) is a pair of prices such that payoffs of the sellers at these prices $(\pi(p, q), \pi(q, p))$ are on the convex hull of the Pareto Frontier and in addition that

$$\pi(p_C, p_C) < \pi(q, p) < \pi(p_M, p_M) < \pi(p, q). \tag{19}$$

16

So, the payoff of one seller is above the monopoly outcome and the other is between the monopoly and competitive outcome. Now consider the following algorithms

$$s^i = \{q \to p, x \to p_C \ for \ x \neq q\}$$
$$s^{-i} = \{p \to q, x \to p_C \ for \ x \neq p\}$$
(20)

This is a grim trigger sort of strategy that aims to sustain the choice of (p, q) in equilibrium. If the discount factor β is sufficiently low these two algorithms form an equilibrium pair (for seller i it is optimal to respond with s^i to s^{-i} and the vice-versa). For sufficiently low β: $\xi(s^i, f^{-i}(s^i)) = (p, q)$ and $\xi(s^{-i}, f^i(s^{-i})) = (q, p)$ in any subgame perfect (and hence Markov perfect) equilibrium. Thus, each seller can guarantee a payoff of at least $\pi(p, q)$ or $\pi(q, p)$ in response to s^{-i} and s^i correspondingly. Finally, since the continuation payoffs $(\pi(p, q), \pi(q, p))$ are on the convex hull of the Pareto frontier, none of the sellers would be able to get a higher payoff by deviating from the current strategy.

Interestingly, such unequal equilibria can also be constructed for arbitrary high discount factors β. If seller $-i$ believes that for any of their response seller i will keep their current strategy s^i, then they also do not have any incentives to deviate from s^{-i}. And for seller i it would indeed be optimal to do that if $\pi(p, q)$ is sufficiently high. To illustrate these ideas, we explicitly construct such an equilibrium with 5 different prices (see Part D of the appendix). As in repeated games, it is however easier to operate in the payoff space and characterize key properties of the set of equilibrium payoffs, which is what we do in this section.

The obvious question to then ask is this: Can we bound the equilibrium payoffs to illustrate the ideas of the inevitability of some collusion and commonality of high collusion? The main result shows that for any Markov equilibrium the payoff of both sellers is bounded from below by the competitive outcome $\pi(p_C, p_C)$, and in any equilibrium at least one seller gets a payoff that is eventually bounded from below by a number that is approximately equal to the monopoly profit $\pi(p_M, p_M)$.

Theorem 2. *Suppose $s_1, s_2, s_3, \ldots$ is a path of algorithms in a Markov equilibrium. Then:*

1. *payoff of both sellers is higher than the competitive outcome:*

$$\forall \ k \geqslant 1, \{u_k^i, v_{k+1}^i\} \geqslant \pi(p_C, p_C) \ \forall \ i \in \{A, B\}, \ and$$

2. *payoff of at least one seller is eventually bounded below approximately by the monopoly outcome:*

$$\exists \ k \ s.t. \ \max_{m \geqslant k}\{u_m^i, v_{m+1}^i\} \geqslant (1 - \beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \pi(p_M, p_M) \ for \ i = A \ or \ i = B.$$

While the first part is straightforward, the second one can be interpreted as follows: If the price grid is dense enough or the effective discount factor is high enough, then in any Markov equilibrium, the payoff of at least one seller is bounded from below by the monopoly profit. Both point to the exercise of some market power by the sellers and taken together they

17

establish this market power is significant. Notice, also, that even for the weakest case of $\beta \approx 0$ and only 3 possible prices in the grid the lower bound it at least $\pi\left(\frac{p_C+p_M}{2}, \frac{p_C+p_M}{2}\right)$ which is non-trivially higher than the competitive payoff.[9]

4.2 Proof of Theorem 2

The proof consists of several parts. In sections 4.2.1-4.2.3 we develop the necessary techniques. In 4.2.1 we show that the state space can be contracted from the set of algorithms to the matrix of price pairs, which significantly reduces the complexity of the analysis. Then in 4.2.2 we derive the conditions that the continuation payoff matrices should satisfy in order for the them to correspond to an equilibrium. Section 4.2.3 then defines a notion of 'best" and "second best" algorithms that each seller could have whenever their competitor is adjusting their algorithm.

Finally, we prove the theorem in section 4.2.4. We consider an arbitrary equilibrium path, and show that the expected continuation payoffs for both sellers tend to increase on this path and if not the corresponding seller can guarantee at worst the "second best" continuation payoff. Lemma 4 then shows that such payoff is above the limit stated in Theorem 2 for at least one of the sellers.

4.2.1 State Space Reduction

Whenever a seller responds with an algorithms that leads to a price pair (p, q) they could choose always the best algorithm consistent with this price pair (that maps q into p). This means that whenever we observe a convergent prices pair (p, q) after seller i adjusted their algorithm, we should observe the same algorithm for seller i. Thus, the price pair (p, q) becomes essentially the state variable, and since there are much fewer prices pairs than possible algorithms this allows us to significantly reduce the state space.

We can define

$$v^i(p, q) = \max_{s: q \to p} v^i(s) \tag{21}$$

$$s^i(p, q) = \arg\max_{s: q \to p} v^i(s) \tag{22}$$

$$u^{-i}(q, p) = u^{-i}(s^i(p, q)) \tag{23}$$

If we know the payoff matrices $v^i(p, q)$ and $v^{-i}(q, p)$ then for any algorithm s of seller $-i$ we can calculate the highest possible continuation payoff that seller i could get:

$$u^i(s) = \max_p (1 - \beta) \cdot \pi(p, s(p)) + \beta \cdot v^i(p, s(p)) \tag{24}$$

Even though this equation does not tell us with which exact algorithms seller i should respond, it pins down the highest payoff they can ensure while responding to an algorithm s by $-i$.

---

[9]The reason why the lower bound is not exactly the monopoly payoff, is that for general payoff matrices with multiple prices there could also exist alternating price equilibria (similar to type III equilibria from Theorem 1). We believe that for specific payoffs functions (e.g. "linear" or "discrete choice") such equilibria might be also ruled out and the lower bound could be improved.

4.2.2 Payoff Matrices in Equilibrium

In this section we will derive the conditions that matrices $v^i(p,q)$ and $u^{-i}(q,p)$ for $i = A, B$ should satisfy if they correspond to payoffs in a certain Markov equilibrium. Also, we will show how to construct the best response to any algorithm if the values of the payoff matrices are known.

First, let us define the worst possible algorithm for seller i to respond to, and the lowest continuation payoff that they can get when responding to some algorithm:

Definition 1. *The worst algorithm of seller $-i$ for seller i is defined as:*

$$\underline{s}^{-i}(p) = \arg\min_q (1 - \beta) \cdot \pi(p,q) + \beta \cdot v^i(p,q) \tag{25}$$

Definition 2. *Then the lowest possible continuation payoff of i when i is adjusting their algorithm is:*

$$\underline{u}^i = \max_p \left[\min_q (1 - \beta) \cdot \pi(p,q) + \beta \cdot v^i(p,q) \right] \tag{26}$$

From equation (24) it follows that for any algorithm s:

$$u^i(s) \geq \underline{u}^i = u^i(\underline{s}^{-i}) \tag{27}$$

Now let's calculate $v^i(p,q)$. Suppose i chooses an algorithm such the next pair of prices (after $-i$ adjusts their algorithm) is (p', q'). Then for this price pair we can define the "before" continuation payoffs $V^i(p', q')$ and $U^{-i}(q', p')$, which include this price pair and assume optimal behavior afterwards:

Definition 3. *"Before" continuation payoffs for price pair (p', q') are defined as*

$$\begin{aligned} V^i(p', q') &= (1 - \beta) \cdot \pi(p', q') + \beta \cdot u^i(p', q') \\ U^{-i}(q', p') &= (1 - \beta) \cdot \pi(q', p') + \beta \cdot v^{-i}(q', p') \end{aligned} \tag{28}$$

Then, if price pair (p', q') follows (p, q) the continuation payoffs are linked as:

$$v^i(p,q) = V^i(p', q') \qquad u^{-i}(q,p) = U^{-i}(q', p') \tag{29}$$

Note that the values of the "before" continuation payoffs $V^i(p', q')$ and $U^{-i}(q', p')$ are directly determined by the "after" continuation payoff matrices $u^i(p', q')$ and $v^{-i}(q', p')$. Finally, the following proposition links the "before" and "after" continuation payoffs, thus establishing the recursive relationship that the continuation payoff matrices $u^i(p', q')$ should satisfy.

Proposition 1. *For any price pair (p, q):*

$$v^i(p,q) = \max V^i(p', q') \qquad s.t. \quad \begin{aligned} &U^{-i}(q', p') \geq \underline{u}^{-i} \\ &U^{-i}(q', p') \geq U^{-i}(q, p) \\ &either\ (p', q') = (p, q)\ or\ q' \neq q \end{aligned} \tag{30}$$

Moreover, if the price pair (p', q') maximizes the above expression then $u^{-i}(q, p) = U^{-i}(q', p')$.

The idea of the proof is as follows. If price pair (p', q') goes after (p, q) then all conditions stated in equation (30) should be satisfied (the payoff of seller $-i$ can never be below $\underline{u}^{-i}$ and it cannot be below $U^{-i}(q, p)$ since $-i$ can keep the strategy from the previous period). In addition (see appendix for the details), by using the following algorithm:

$$s_*^i(p, q) = \{q \to p, q' \to p', x \to \underline{s}^i(x) \: for \: x \neq q \: and \: x \neq q'\} \tag{31}$$

seller i can make sure that the subsequent prices will actually be (p', q').

Proposition 1 gives us a powerful tool showing which price transitions are possible in an equilibrium. In addition, it gives us precise conditions that payoff matrices $v^i(p, q)$ and $u^{-i}(q, p)$ should satisfy in order for them to correspond to an equilibrium.

4.2.3 "First best" and "Second best" Continuation Payoffs

The final piece of notation that we would need in order to prove Theorem 2 is the first best and the second best continuation payoffs and the corresponding algorithms that would allow sellers to achieve such payoffs.

Definition 4. *The "best" continuation payoff of i when $-i$ is adjusting their algorithm is defined as:*

$$v_*^i = \max V^i(p, q) \qquad s.t. \: U^{-i}(q, p) \geq \underline{u}^{-i} \tag{32}$$

Also if (p_*^i, q_*^i) maximize the above expression then we define the best continuation algorithm of i as [10]:

$$s_*^i = \{q_*^i \to p_*^i, x \to \underline{s}^i(x) \: for \: x \neq q_*^i\} \tag{33}$$

Definition 5. *If (p_*^i, q_*^i) are the "first best" prices from definition 4 then we define the "second best" continuation payoff of seller i when $-i$ is adjusting their algorithm as:*

$$v_{**}^i = \max V^i(p, q) \qquad s.t. \: \begin{array}{c} U^{-i}(q, p) \geq \underline{u}^{-i} \\ q \neq q_*^i \end{array} \tag{34}$$

Also if (p_{**}^i, q_{**}^i) maximizes the above expression (34) then we can define the "second best" continuation algorithm of i as:

$$s_{**}^i = \{q_{**}^i \to p_{**}^i, x \to \underline{s}^i(x) \: for \: x \neq q_{**}^i\} \tag{35}$$

Note that seller i cannot get a continuation payoff $v^i(s)$ higher than v_*^i for any s. Moreover, the payoff v_*^i is exactly achievable for $s = s_*^i$ (if seller $-i$ does not continue with prices (p_*^i, q_*^i) their payoff is at most $\underline{u}^{-i}$). The second best payoff v_{**}^i corresponds to the case when seller i cannot use the opponent's price q_*^i from the first best pair (p_*^i, q_*^i).[11]

---

[10] In Definitions 4 and 5 if several feasible price pairs maximize $V^i(p, q)$ then in accordance with assumption 3 from Section 2.3 we select the one maximizing $U^{-i}(q, p)$.

[11] If none of the price pairs satisfy the stated conditions for the "second-best" (this means seller $-i$ always

4.2.4 Evolution of Equilibrium Continuation Payoffs

Consider some arbitrary equilibrium sequence of algorithms $s_0, s_1, s_2, \ldots$ Note that "even" algorithms in the sequence $s_0, s_2, \ldots$ correspond to one seller, and "odd" algorithms to the other. Below when we write expressions like $u^i(s_k)$ we assume that i and k agree with each other (e.g. in this case it means that s_k is the algorithm of seller i).

Lemma 1 shows that each seller can guarantee that their the continuation payoffs stay above the competitive payoff $\pi(p_C, p_C)$ which proves the first part of Theorem 2.

Lemma 1. $u^i(s_k) \geq \pi(p_C, p_C)$ *for* $k \geq 0$ *and* $v^{-i}(s_k) \geq \pi(p_C, p_C)$ *for* $k \geq 1$.

We can also show that the continuation payoffs tend to increase on the equilibrium path (and if not the corresponding seller can guarantee at least the second best continuation payoff).

Lemma 2. *For any k and corresponding i:*

1. $u^{-i}(s_{k+1}) \geq v^{-i}(s_k)$

2. *Either $v^i(s_{k+1}) \geq u^i(s_k)$ or $v^i(s_{k+1}) = v^i_*$ or $v^i(s_{k+1}) = v^i_{**}$. Moreover, the last case is only possible if $\xi(s_k, s_{k+1}) = (q^i_*, x)$ for some x where q^i_* is from definition 4.*

Suppose the continuation payoffs of one seller becomes constant, then we can show that the continuation payoffs of the other seller will also become constant. Moreover, at least on of these constant payoffs should be above the monopoly profit.

Lemma 3. *If for some i and k:*

$$u^i(s_k) = v^i(s_{k+1}) = u^i(s_{k+2}) = \ldots \tag{36}$$

then for some m

$$u^{-i}(s_m) = v^{-i}(s_{m+1}) = u^{-i}(s_{m+2}) = \ldots \tag{37}$$

Moreover, either $u^i(s_k) \geq \pi(p_M, p_M)$ or $u^{-i}(s_m) \geq \pi(p_M, p_M)$.

Finally, if continuation payoffs for both sellers never become constant then eventually both of them will be above the corresponding second best payoffs v^i_{**}. Thus, to prove theorem 2 we only need to show that at least one of these second best payoffs is above the bound stated in the theorem.

Lemma 4.

$$\max\left(v^i_{**}, v^{-i}_{**}\right) \geq (1 - \beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \pi(p_M, p_M) \tag{38}$$

ends up with price q_* after adjusting their algorithm), it is fairly easy to show that the continuation payoffs for at least one of the sellers in any equilibrium sequence will exceed $\pi(p_M, p_M)$. See part B of the appendix for details. So, in what follows we will assume that the second-best always exists.

5 Extensions and robustness considerations

The main message of this section is that if the reader is worried how the main results would extend to scenarios where Assumptions 1 and 2 are relaxed; well, they shouldn't be too worried. If the reader ponders what would happen if we extend the equilibrium notion from Markov perfect to subgame perfect; the sharpness of results is tempered by the multitude of equilibria and quite complex strategies are required to sustain these.

5.1 Alternative Ways to Resolve Cycles

For our main results we assumed that the sellers cannot respond with algorithms that generate cycles (Assumption 2 in Section 2.3). This is essentially a restriction on the strategy space. In this section we relax this assumption for the two prices model, presented in Section 3. Recollect from Table 4 that whenever s_T is chosen by one seller and s_R by the other, prices go into a cycle. We consider two ways of resolving cycles:

1. A cycle is equivalent to having the lowest prices from each seller in the cycle. This is motivated by the fact that if prices change frequently then customers can learn about this and time their purchase so that they can get the lowest possible price:

$$\pi(s_T, s_R) = \pi(s_R, s_T) = \pi(p_C, p_C) = 0 \tag{39}$$

2. A seller's payoff from a cycle is equivalent to the average payoff from all price pairs that are observed in the cycle. This means customers arrive continuously and do not optimize on timing:

$$\pi(s_T, s_R) = \pi(s_R, s_T) \quad \begin{aligned} &= \frac{\pi(p_C, p_C) + \pi(p_M, p_C) + \pi(p_M, p_M) + \pi(p_C, p_M)}{4} \\ &= \frac{2 + x - y}{4} < \frac{3}{4} \end{aligned} \tag{40}$$

We can first show that for both options it is never optimal to respond to s_T with s_R.

Lemma 5. $f^i(s_T) \neq s_R \ \forall \ i \in \{A, B\}$.

Proof. See appendix. $\square$

Lemma 5 implies that the first half of the proof of Theorem 1 is unchanged, which means the best response to s_C can still only be s_T and after s_T we still observe only monopoly prices. Now it is not difficult to describe all possible equilibria with cycles.

Proposition 2.

1. *If cycles are equivalent to the lowest price (case 1) then all equilibria are exactly the same same as described in Theorem 1*

2. *If cycles are equivalent to the average payoff (case 2) then*

(a) for $x \leq \beta$ and $y \geq 4\beta - 2 - 3x$ the following Markov equilibria are possible

$$\begin{aligned}
s_R &\to s_C \to s_T \\
s_T &\to \{s_T, s_M\} \qquad (type\,I) \\
s_M &\to \{s_T, s_M\}
\end{aligned} \qquad (41)$$

(b) for $x \leq \beta$ and $y \leq 4\beta - 2 - 3x$ the following Markov equilibria are possible

$$\begin{aligned}
s_R &\to s_T \\
s_T &\to \{s_T, s_M\} \qquad (type\,I') \\
s_M &\to \{s_T, s_M\}
\end{aligned} \qquad (42)$$

(c) for $x > \beta$ the equilibria are the same as described in Theorem 1.

Thus, the only difference from Theorem 1 is that in case 2 for some values of parameters it might be optimal to switch directly to s_T from s_R without the intermediate step of s_C in an equilibrium of type I. Type II and III equilibria, where the latter can lead to alternating price sequences, exist for exactly the same values of parameters x and y. And, as before, for majority of the parameters monopoly price is the unique convergent outcome in equilibrium.

5.2 Precise Calculations of Expected Payoffs

In this extension we will relax Assumption 1 from Section 2.3 and show that our results generally hold (for sufficiently small interval between price adjustments Δt) even if we calculate payoffs precisely without ignoring initial convergence and discreteness of price adjustments. First we rewrite the precise model incorporating pre-convergent prices and calculate expected payoffs within this model. Then we study the Markov equilibria of this updated framework.

Since the updated model and corresponding results require the introduce of new notation, we relegate the formalization to Part B of the appendix. The main insight can be (informally) stated thus:

Claim. *For sufficiently small Δt, every Markov equilibrium in the "precise model" generically corresponds a Markov equilibrium of the "main model" we study by resolving the dependence on the initial prices arbitrarily.*

5.3 Subgame Perfect Equilibria

In this extension we consider the more general subgame perfect equilibria (SPE) and try to give a sense of the possible outcomes. We restrict our attention to the simplest two price case. The sellers still adjust their algorithms asynchronously, however, now the response of each might depend not only on the current algorithm of the opponent, but on the overall history of algorithms up to this point.

Suppose one of the sellers is adjusting their algorithm and suppose the current algorithm of their opponent is s. Define the set of all possible continuation payoffs (u, v) in a subgame

perfect equilibrium depending on s (u corresponds to the seller adjusting their algorithm and v to the other one):

$$H(s) = \{(u,v) : (u,v) \text{ is possible in SPE starting with } s\} \tag{43}$$

There is a recursive way to construct these sets in the spirit of Abreu, Pearce, and Stacchetti [1990]. Suppose the seller adjusting their algorithm decides to respond to s with s', then the set of possible continuation payoffs is as follows:

$$H(s,s') = \left\{(u,v) : \begin{array}{l} u = (1-\beta) \cdot \pi(s',s) + \beta \cdot v' \\ v = (1-\beta) \cdot \bar{\pi}(s,s') + \beta \cdot u' \\ (u',v') \in H(s') \end{array}\right\} \tag{44}$$

The lowest possible continuation payoff u of the seller adjusting their algorithms is then:

$$u_{min}(s) = \max_{s'} \min_{(u,v) \in H(s,s')} u \tag{45}$$

Finally the set of possible SPE payoffs starting with s is:

$$H(s) = F(H)(s) = \bigcup_{s'} \left\{(u,v) : \begin{array}{l} (u,v) \in H(s,s') \\ u \geq u_{min}(s) \end{array}\right\} \tag{46}$$

Any (u,w) with $u \geq u_{min}(s)$ could be sustained in an SPE since for any deviation their exists a subgame which generates a payoff not higher than u.

Now to construct the sets of possible continuation payoffs $H(s)$ we can start with some arbitrary payoff sets that are guaranteed to contain $H(s)$, for example (here $\pi_{min} = \min_{p,q} \pi(p,q)$ and $\pi_{max} = \max_{p,q} \pi(p,q)$):

$$H_0(s) = [\pi_{min}, \pi_{max}] \times [\pi_{min}, \pi_{max}] \tag{47}$$

And, to calculate the sets $H(s)$ we can use the mapping F from equation (46) (which maps the set of payoffs $H = \{H(s) \text{ for all } s\}$ onto itself) and then construct the sets of possible SPE payoffs as:

$$\begin{array}{l} H_{n+1} = F(H_n) \\ H(s) = \lim_{n\to\infty} H_n(s) \end{array} \tag{48}$$

Note that for any n: $H_{n+1}(s) \subset H_n(s)$, thus, the limit exists. Moreover, the limit is exactly $H(s)$: for any $(u,v) \in \lim_{n\to\infty} H_n(s)$ we can construct a SPE by recursively calculating the possible continuation payoffs in the next periods (which always stay within $\lim_{n\to\infty} H_n(s)$ and thus can always be extended further).

To illustrate the possible equilibria consider the two price setup with payoffs for each customer from Table 5. For given payoffs we can construct the the sets $H(s)$ numerically[12].

---

[12]We do that by splitting the initial set $H_0(s)$ into a fine grid of squares (10000 by 10000). Then for each n
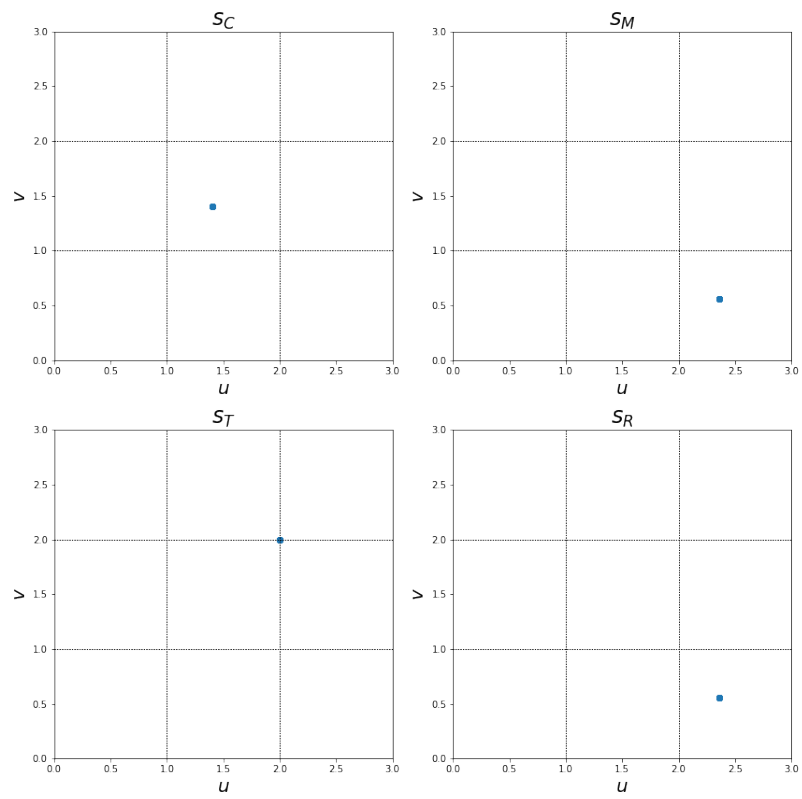
(a) $\beta = 0.4$
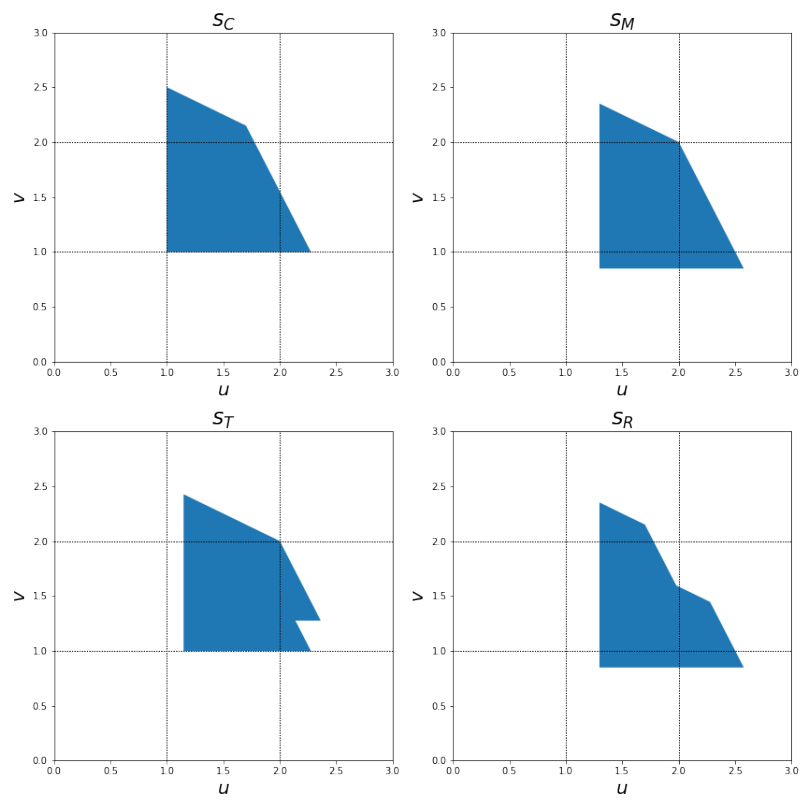
(b) $\beta = 0.85$

Figure 3: Possible continuation payoffs in a subgame perfect equilibrium.

$A \backslash B$	p_M	p_C
p_M	2, 2	0, 3
p_C	3, 0	1, 1

Table 5: Payoffs for SPE example.

Figure 3 show the sets of possible payoffs for each algorithms of the opponent for two values of the discount factor β. We make two points through these simulations.

Pure monopoly for low discounting. If the discount factor is low, then the only possible SPE leads to monopoly outcomes (p_M, p_M). It is always optimal to respond with s_T to s_T or s_C. This confirms the result asserted in Salcedo [2015] that as the effective discount factor converges to zero, repeated play of the monopoly outcome emerges as the unique SPE.

In fact our simulations suggest that the threshold on effective discount factor doesn't have to be in the neighborhood of zero. As can be seen in Figure 3, this holds for $\beta = 0.4$ for any starting strategy. The theoretically interesting fact is that in a standard repeated games setting, low discounting would have led to the uniqueness of (p_C, p_C) as the equilibrium outcome.

High discounting leads to large and complex equilibrium payoff sets. As is widely understood from folk theorems, for high discount factors equilibrium expands to maximal possible size. This continue to hold for the asynchronous repeated we study. Figure 3 shows the set of equilibrium payoffs for $\beta = 0.85$. Note that both competition, i.e. payoff of $(1,1)$, and conflict, i.e. payoff even lower than 1 for one of the sellers, can be sustained. The strategies required to sustain these payoffs though are quite complex.

Consider for example the competitive equilibrium in which each seller always chooses s_C which would generate continuation payoffs of $(1, 1)$ for the payoff matrix from Table 5. Suppose one of the sellers deviates to s_T in one of the rounds, and subsequent continuation payoffs are $(u, v) \in H(s_T)$. In order for this deviation to be not profitable, v should be equal to 1. But note that u is non-trivially different from 1 $(u \geq 1.15)$, thus, we should expect a complicated sequence of price pairs and corresponding algorithms that would generate such payoffs. For example, consider the lowest point $(u = 1.15$ and $v = 1)$. For this point we can construct the possible continuation sequences of algorithms that could generate such payoffs. We do this by choosing the algorithm s' such that $(u, v) \in H(s_T, s')$ and then doing this recursively. In the end, we get the following sequence of algorithms ("$\times 13$" means s_C is repeated 13 times):

$$s_T, s_C \times 13, s_M, s_C, s_M, s_C \times 2, s_M \times 10, \dots, {}^{13} \qquad (49)$$

the sets H_{n+1} contain only the squares for which at least one point (u, v) is a feasible with continuation payoff from H_n. We also add a small margin allowing for mistakes of numerical calculations. Thus, the constructed sets are guaranteed to include $H(s)$ though could be a bit larger.

[13]Here the possible continuation algorithms were always reviewed in the following order: s_C, s_M, s_T, s_R. Thus, if there were several possibilities with which algorithm could continue the earliest one in the above sequence was always selected. This also means that other sequences of algorithms are certainly possible. However, all such sequences have similar irregular structure (we also constructed all possible sequences up to 15 elements and manually checked them).

The above sequence of responses requires some very specific beliefs about how the opponent should behave in equilibrium while choosing algorithms, which we find unrealistic. This is the reason why we decided to focus only on Markov perfect equilibria, which restrict the complexity of the possible strategies.[14]

6 A comparison to Salcedo [2015]

The above extension to subgame perfect equilibrium provides a good segue to compare our analysis with Salcedo [2015]. They also look at a duopoly model of price-setting with probabilistic demand and algorithms as strategies, but use subgame perfectness as the equilibrium criterion. The basic idea of using such a set up to explore algorithmic collusion is elegant. We depart from their analysis in several ways.

First, we model the competition in algorithms differently. In Salcedo [2015], whenever a customer arrives, the price is set by a complex automaton. In our model customer arrivals are unobservable and the algorithms react only to prices set by the competitor. We believe this assumption lends a realism in how algorithms actually interact, especially when we think about online sellers. In addition, it allows us to separate the frequency of customer arrivals from the frequency of possible prices changes.

Further, we assume that algorithms can react fast to possible prices changes by the competitor, again a reasonable assumption. This allows us to replace the complicated continuous time game with a simpler a asynchronous repeated game in which payoffs in each period are determined by the convergent prices for the current algorithms. This, in turn, allows to derive sharper results, which would be much more difficult to obtain otherwise.

The second key difference is that in Salcedo [2015] the main result about inevitability of monopoly outcomes is stated only for the limiting case in which the frequency of algorithmic adjustments (μ in our model) converges to 0.[15] Then the effective discount factor $\beta = \frac{\mu}{\mu+r}$ also converges to 0. This limits the dynamic dimension of the strategic interaction between sellers, for they essentially behave "myopically," best responding only to the current algorithm of the opponent, not taking into account how the opponent would react to their choice.

The case of $\beta \approx 0$ is also unrealistic. Suppose we fix the discount factor at $r = 5\%$ per year. Then to generate $\beta = 0.5$, which is not even that low, we would need $\mu = 0.05$. In plain speak, this means the sellers have the opportunity to revise their algorithms, on average, once in 20 years. If $\beta = 0.1$, then it is once in 180 years. More realistically, suppose we fix $\mu = 5$, which means each seller has on average five opportunities per year to revise their algorithm. Then, $\beta \approx 0.99$. For $\mu = 10$, $\beta \approx 0.995$.

When β is high, as we show in Section 5.3, the monopoly outcome is no longer inevitable

---

[14]In terms of the repeated games literature if we allowed for public randomization, then a lot of the complexity in achieving certain payoffs can be pushed into the randomization devices. It is not clear, however, how to think about public randomization when it come to modeling algorithms.

[15]Salcedo [2015] also states a result for revision opportunity going to infinity. This means that sellers are changing algorithms faster than consumer arrivals, so before any payoffs accrue, and in the limit even faster than price changes, which is unrealistic.

in subgame perfect equilibria, and a wide range of outcomes are possible (as is typical for repeated games). Nevertheless, we are able to show that some collusion is inevitable and extreme collusion is very likely for Markov perfect equilibria. Thus, we can conclude that even for high β's, which we argue are more realistic, algorithms will still facilitate collusion.

7 Final remarks

The ancient treatise on statecraft, Arthshastra (written and then chiseled between 3rd century BCE and 2nd centruy AD), accredited in folklore to a writer Kautilya, makes several observations about the act of what we term collusive practices, their harm to consumers, and necessity of regulatory practices to reign them in:

> [T]raders, joining together and raising or lowering the (prices of) goods, make a profit of one hundred *panas* on one *pana* or of hundred *kumbhas* on one *kumbh*a. (VIII.4.36, emphasis added) (Pana, a silver coin, was the basic unit of money, whereas kumbha refers to a measure of weight (II.29.32))

> For traders who by *conspiring* together hold back wares or sell them at a high price, the fine is one thousand *panas*. (IV.2.19, emphasis added)[16]

Closer to modern times, the Sherman Act of 1890 was the the first federal legislation in the United States to put down the promotion of free competition an a desired objective of the state. Organized societies has always faced the possibility of extraction of consumer surplus by cartels of producers beyond what is deemed legitimate. But, this line of legitimacy and the means used to cross it without incurring the wrath of the law evolve over time. In the twenty first century, rapid rise in computational technology has arguably equipped firms with a means to collude (almost inadvertently) that is outside the scope of the current antitrust cannon—by delegating pricing to algorithms.

Algorithmic pricing and its various implications are now a rapidly burgeoning field of enquiry in economics. The purpose of this paper is to offer a simple enough theoretical model that shows why the quickness of response and temporary commitment offered by algorithms lead to inevitability of some collusion and commonality of extreme collusion. The intuition is communicated sharply first through a two-price model and then a generalization to multiple prices shows the basic insights carry through.

Technically speaking, the analysis here is restricted to two firms and Markov perfect equilibrium. Generalizing to multiple firms and subgame perfect equilibrium with some selection criterion seems like a worthwhile exercise. We also restricted the algorithm to be a function of the current price of the opponent. It could potentially depend on a richer history. In the language of automaton, if the number of internal states of algorithms is not very large (eg. incorporating strategies like grim trigger), then we conjecture that the main results should still go through. This is because even though the initial converge period could be longer, but if

---

[16]These are cited from Kumar [2012].

Δt is small enough we can similarly focus on the final price pair. We also discuss in Appendix E why explicitly considering experimentation is easy to handle in our model, since the cost is linear in the set of prices and in unit of time, and hence small. However, if we allow for the aforementioned richer class of algorithms then these experimentation costs can become non-trivial and tackling this trade-off is also an interesting question for future work.

Finally, algorithmic pricing and information technology more generally is now widely recognized to be a challenge for anti-trust regulation. How to incorporate these issues into models of mergers and acquisitions is a fascinating topic for future research.

8 Appendix

The appendix is divided into four parts. Part A contains the proofs missing in the main text. Part B shows that Theorem 2 continues to hold even if the "second best" price pair from definition 5 may not exist, in fact in that case the results is strengthened. Part C presents in detail the extension introduced in Section 5.2—accounting for precise calculations of payoffs which includes trades before prices converge. Part D constructs an unequal equilibrium for an example with five possible prices. Part E discusses the finer details of explicitly modeling experimentation in our framework.

Part A: Proofs

Proof of Theorem 1.

1. For any seller i and any s: $\xi(s, f^i(s)) \neq (p_C, p_M)$. Otherwise seller $-i$ can keep the strategy s, generating the highest possible payoff $\pi(p_C, p_M)$ for them and the lowest possible $\pi(p_M, p_C)$ for their competitor, which cannot be an equilibrium. Thus, the best response to s_C can only be s_C or s_T and the the best response to s_R cannot be s_M.

2. $f^i(s_T) \neq s_C$.

   - Assume the contrary, then the following equilibrium sequences $s_0, s_1, s_2, \ldots$ starting with s_T are possible:
   $$
   \begin{aligned}
   &s_T, s_C, s_T, s_C, \ldots \\
   &s_T, s_C, s_C, s_C, \ldots \\
   &s_T, s_C, s_C, s_T, \ldots
   \end{aligned}
   \tag{50}
   $$

   In each case seller i would find it optimal to deviate to $s_1 = s_T$.

3. Since $f^i(s_T) \in \{s_T, s_M\}$, we have $\xi(s_T, f^i(s_T)) = (p_M, p_M)$. Moreover, after we observe a price pair (p_M, p_M), all subsequent prices are (p_M, p_M) as well (after (p_M, p_M) each could use s_T and guarantee the discounted payoff of at least (p_M, p_M) for them).

4. $f^i(s_C) = s_T$.

29

- Assume the contrary, then $f^i(s_C) = s_C$. Then the following equilibrium sequences $s_0, s_1, s_2, \ldots$ starting with s_C are possible

$$
\begin{aligned}
s_C, s_C, s_C, s_C, \ldots \\
s_C, s_C, s_T, \ldots
\end{aligned}
\tag{51}
$$

  In each case seller i would find it optimal to deviate to $s_1 = s_T$.

5. If $f^{-i}(s_R) = s_C$ then $f^i(s_R) = s_C$.

- Assume the contrary, then $f^i(s_R) = s_R$ and we will observe the following equilibrium sequence starting with s_R for seller i:

$$
s_R, s_R, s_C, s_T, \ldots
\tag{52}
$$

  For such sequence seller i would find it optimal to deviate to $s_1 = s_C$.

6. $f^i(s_R) = s_R$ is possible only if $y < \beta \cdot (x - \beta)$

- Suppose $f^i(s_R) = s_R$ then $f^{-i}(s_R) = s_R$ as well.
- The expected payoff from responding with s_R if your competitor's algorithm is s_R:

$$
(1-\beta)\cdot[\pi(p_C,p_M) + \beta \cdot \pi(p_M,p_C) + \ldots] = (1-\beta)\cdot\pi(p_C,p_M) + \beta\cdot\frac{\pi(p_M,p_C) + \pi(p_M,p_C)}{1 + \beta}
\tag{53}
$$

  The expected payoff from responding with s_C if your competitor's algorithm is s_R (the corresponding equilibrium sequence would be $s_R, s_C, s_T, \ldots$):

$$
(1 - \beta) \cdot \pi(p_C,p_M) + \beta(1 - \beta) \cdot \pi(p_C,p_C) + \beta^2 \cdot \pi(p_M,p_M)
\tag{54}
$$

  Thus, the response s_R is optimal when

$$
\frac{\pi(p_M,p_C) + \pi(p_M,p_C)}{1 + \beta} > (1 - \beta) \cdot \pi(p_C,p_C) + \beta \cdot \pi(p_M,p_M)
\tag{55}
$$

  or equivalently when $y < \beta \cdot (x - \beta)$

7. $f^i(s_M) = s_C$ if $x < \beta$. Otherwise $f^i(s_M) \in \{s_M, s_T\}$.

- The expected payoff from responding with s_C if your competitor's algorithm is s_M (the corresponding equilibrium sequence would be $s_M, s_C, s_T, \ldots$):

$$
(1 - \beta) \cdot \pi(p_C,p_M) + \beta(1 - \beta) \cdot \pi(p_C,p_C) + \beta^2 \cdot \pi(p_M,p_M)
\tag{56}
$$

- The expected payoff from responding with s_T if your competitor's algorithm is s_M (the corresponding equilibrium sequence would be $s_M, s_T, \ldots$):

$$
(1 - \beta) \cdot \pi(p_M,p_M) + \beta(1 - \beta) \cdot \pi(p_M,p_M) + \beta^2 \cdot \pi(p_M,p_M)
\tag{57}
$$

30

The first option is preferable if

$$\pi(p_C, p_M) + \beta \cdot \pi(p_C, p_C) > (1 + \beta) \cdot \pi(p_M, p_M) \tag{58}$$

or equivalently if $x > \beta$

- Note, also, that by responding with s_M to s_M you cannot get a higher expected payoff than by responding with s_T (and you might only do that the expected payoffs are identical).

These points all together prove the result.

Proof of Corollary 1

We can directly verify that in type I or II equilibria we always end up with monopoly prices. In type III equilibrium if we start with s_M or s_R then the subsequent responses by each seller will always be s_R, which would generate the corresponding alternating price sequence. Type III equilibria are only possible if $x > \beta$ and $y < \beta \cdot (x - \beta)$, which is equivalent to the condition stated in the corollary.

Proof of Corollary 2

1. The first statement directly follows the definitions of x and y from equation (16).

2. To prove the sufficient condition for "linear" demand we can simply calculate the following expression:

$$[\pi(p_M, p_M) - \pi(p_M, p_M - x)] + [\pi(p_M - x, p_M - x) - \pi(p_M - x, p_M)] = \\ p_M \cdot \frac{\alpha \cdot x}{2D} + (p_M - x)\frac{-\alpha \cdot x}{2D} = \frac{\alpha x^2}{2D} > 0 \tag{59}$$

3. Now consider the "discrete choice" profit function. Without loss of generality we can assume $b = 1$ (we could always scale prices so that $b = 1$) and replace a with $a \cdot e^{-p_M}$:

$$\pi(p, q) = \frac{p \cdot e^{-(p - p_M)}}{a + e^{-(p - p_M)} + e^{-(q - p_M)}} \tag{60}$$

Then

$$\pi(p_M, p_M) + \pi(p_M - x, p_M - x) > \pi(p_M, p_M - x) + \pi(p_M - x, p_M) \qquad \Leftrightarrow \tag{61}$$

$$\frac{p_M}{a + 2} + \frac{(p_M - x)e^{-x}}{a + 2e^{-x}} > \frac{p_M + (p_M - x)e^{-x}}{a + 1 + e^{-x}} \tag{62}$$

Finally since for any positive A, B, C, D:

$$\frac{A}{B} + \frac{C}{D} > 2\frac{A + C}{B + D} \qquad \Leftrightarrow \qquad (B - D) \cdot \left(\frac{A}{B} - \frac{C}{D}\right) > 0 \tag{63}$$

31

the inequality (62) is equivalent to

$$2(1 - e^{-x}) \cdot \left(\frac{p_M}{a+2} - \frac{(p_M - x)e^x}{a + 2e^{-x}} \right) > 0 \tag{64}$$

which holds since p_M is the monopoly price.

Proof of Proposition 1

1. First, let us show that all the conditions are necessary for any price pair (p', q') that follows after (p, q). (This implies that $v^i(p, q)$ cannot be above the maximum.)

   - Any algorithm s consistent with (p, q) should map $q \to p$ which gives us the third condition.

   - In addition, when seller $-i$ gets the opportunity to revise their algorithm, they can always keep the pair (p, q) (e.g. by not changing their previous algorithms) and thus $U^{-i}(q', p')$ cannot be below than $U^{-i}(q, p)$, which gives the second condition.

   - Finally, since for any s: $u^{-i}(s) \geq \underline{u}^{-i}$ then first condition should also also hold.

2. Now suppose (p', q') maximizes $V^i(p', q')$ and satisfies all the necessary conditions. Then i could use the following algorithm

$$s_*^i(p, q) = \{q \to p, q' \to p', x \to \underline{s}^i(x) \, for \, x \neq q \, and \, x \neq q'\} \tag{65}$$

   to which $-i$ would respond with some s^{-i} such that $\xi(s_*^i(p, q), s^{-i}) = (p', q')$ with the continuation payoff $u^{-i}(s_*^i(p, q)) = U^{-i}(q', p')$. Otherwise:

   - If $\xi(s_*^i(p, q), s^{-i}) = (p, q)$ then $u^{-i}(s_*^i(p, q)) \leq U^{-i}(q, p) \leq U^{-i}(q', p')$
   - If $\xi(s_*^i(p, q), s^{-i}) \neq (p, q) \, or \, (p', q')$ then $u^{-i}(s_*^i(p, q)) \leq \underline{u}^{-i} \leq U^{-i}(q', p')$

Proof of Lemma 1

1. Each seller when adjusting their algorithm always has an option of choosing the competitive price p_C and keep it forever. Thus $u^i(s_k) \geq \pi(p_C, p_C)$.

2. Consider some $k \geq 1$ suppose $\xi(s_{k-1}, s_k) = (q, p)$. Then

$$u^{-i}(s_{k-1}) = (1 - \beta) \cdot \pi(p, q) + \beta \cdot v^{-i}(s_k) \geq \pi(p_C, p_C)$$

3. If $\pi(p, q) \leq \pi(p_C, p_C)$ then $v^{-i}(s_k) \geq \pi(p_C, p_C)$. If $\pi(p, q) > \pi(p_C, p_C)$ then seller $-i$ can choose $s_k' = \{q \to p, x \to x \, for \, x \neq q\}$ and keep it forever, with the payoff of at least $\pi(p_C, p_C)$. Thus, $v^{-i}(s_k) \geq \pi(p_C, p_C)$.

32

Proof of Lemma 2

1. Note that $v^{-i}(s_k) = (1-\beta) \cdot \bar{\pi}(s_k, s_{k+1}) + \beta \cdot u^{-i}(s_{k+1})$. Then, seller $-i$ can keep the strategy s_k and generate the continuation payoff of at least $\bar{\pi}(s_k, s_{k+1})$, which proves the first statement of the lemma.

2. Note that $u^i(s_k) = (1-\beta) \cdot \pi(s_{k+1}, s_k) + \beta \cdot v^i(s_{k+1})$. Suppose $\xi(s_k, s_{k+1}) = (q, p)$.

3. If $q \neq q_*^i$ then seller i can use the following algorithm[17]

$$s^i = \{s_*^i | q \to p\} \tag{67}$$

the continuation payoff for which is either v_*^i or not lower than $\pi(p, q) = \pi(s_{k+1}, s_k)$.

4. If $q = q_*^i$ then seller i can use the following algorithm

$$s^i = \{s_{**}^i | q \to p\} \tag{68}$$

the continuation payoff for which is either v_{**}^i or not lower than $\pi(p, q) = \pi(s_{k+1}, s_k)$.

Proof of Lemma 3

1. Suppose $u^i(s_k) = v^i(s_{k+1}) = u^i(s_{k+2}) = \cdots = c$. This means that each individual payoff of seller i is equal to c:

$$\pi(s_{k+1}, s_k) = \bar{\pi}(s_{k+1}, s_{k+2}) = \pi(s_{k+3}, s_{k+2}) = \cdots = c \tag{69}$$

2. Now consider the highest individual payoff of seller $-i$ in the sequence:

$$\bar{\pi}(s_k, s_{k+1}), \ \pi(s_{k+2}, s_{k+1}), \ \bar{\pi}(s_{k+2}, s_{k+3}), \ \ldots \tag{70}$$

Then the corresponding seller can keep the strategy, in response to which this highest payoff was generated. This would lead to the same continuation payoff c for seller i and the highest possible constant continuation payoff for $-i$.

3. Assume the contrary. This implies that $\underline{u}^{-i} < \pi(p_M, p_M)$ and $\underline{u}^i < \pi(p_M, p_M)$. This together with Lemma A2 implies that an equilibrium with monopoly prices is possible and after (p_M, p_M) all the subsequent prices are (p_M, p_M) as well.

4. Suppose for some k: $\xi(s_k, s_{k+1}) \neq (p_M, x)$ for some $x < p_M$ then instead of s_{k+1} the corresponding seller could use the following strategy

$$s'_{k+1} = \{s_{k+1} | p_M \to p_M\} \tag{71}$$

---

[17]Here and below we use the following notation to create a new algorithm s' by slightly adjusting the existing one s:

$$s' = \{s | x \to y\} \qquad \Leftrightarrow \qquad s'(p) = \begin{cases} s(p) & if \ p \neq x \\ y & if \ p = x \end{cases} \tag{66}$$

33

which would lead to the same current payoff and monopoly equilibrium afterwards.

5. If the above k does not exist then for any k for some $x, y < p_M$:

$$\xi(s_k, s_{k+1}) = (p_M, x) \qquad \xi(s_{k+1}, s_{k+2}) = (p_M, y) \tag{72}$$

Then the corresponding seller instead of s_{k+2} can use the following strategy

$$s'_{k+2} = \{s_{k+2} | p_M \to p_M\} \tag{73}$$

which would generate the current payoff of at least

$$\pi(s'_{k+2}, s_{k+1}) \geq \pi(x, p_M) = \pi(p_M, y) \tag{74}$$

and again the subsequent monopoly equilibrium.

Proof of Lemma 4

We prove the lemma by considering possible cases in terms of the first-best prices pairs (p^i_*, q^i_*) from Definition 4. Since we need both prices pairs at the same time to simplify exposition we use the following notation:

$$
\begin{aligned}
(p_*, q_*) &= \arg\max V^i(p, q) & s.t.\ U^{-i}(q, p) \geq \underline{u}^{-i} \\
(\dot{q}, \dot{p}) &= \arg\max V^{-i}(q, p) & s.t.\ U^i(p, q) \geq \underline{u}^i
\end{aligned}
\tag{75}
$$

Similarly, we will use (p_{**}, q_{**}) and $(\ddot{q}, \ddot{p})$ as the prices corresponding to the second-best maxima from Definition 5. In general, in this section all prices indexed with letter p (except for $p_C, p_{C+\Delta}, \ldots, p_{M-\Delta}, p_M$) would correspond to seller i and indexed with letter q would correspond to seller $-i$.

For the proof we also need Lemmas A1, A2, A3 stated below and proven after. Note that, if $U^{-i}(q, p) \geq \underline{u}^{-i}$ then seller i can use an algorithm $s = \{q \to p, x \to \underline{s}^i(x)\ for\ x \neq q\}$ to generate the expected payoff of at least $\pi(p, q)$, and, thus:

$$
\begin{aligned}
v^i(p, q) &\geq v^i(s) \geq V^i(p, q) \geq \pi(p, q) \\
u^i(p, q) &\geq U^i(p, q) = (1 - \beta) \cdot \pi(p, q) + \beta \cdot v^i(p, q) \geq \pi(p, q)
\end{aligned}
\tag{76}
$$

Similar inequalities holds for seller $-i$ if $U^i(p, q) \geq \underline{u}^i$.

Lemma A 1. *Suppose $\pi(p, q) \geq \pi(p_C, p_C)$ and $\pi(q, p) \geq \pi(p_C, p_C)$ and either $(p, q) \neq (\dot{p}, q_*)$ then either $U^i(p, q) \geq \underline{u}^i$ or $U^{-i}(q, p) \geq \underline{u}^{-i}$ (or both).*

Lemma A 2. *Either $U^i(p_M, p_M) \geq \underline{u}^i$ or $U^{-i}(p_M, p_M) \geq \underline{u}^{-i}$ (or both).*

Lemma A 3. *If $q_* = p_M$ then for any $p \neq p_*$: either $\pi(p_*, q_*) > \pi(p, q_*)$ or $\pi(q_*, p_*) > \pi(q_*, p)$. Moreover, $V^i(p_*, q_*) = \pi(p_*, q_*)$ and $U^{-i}(q_*, p_*) = \pi(q_*, p_*)$.*

34

Case $q_* \neq p_M$ and $\dot{p} \neq p_M$

From Lemma A2 either $U^i(p_M, p_M) \geq \underline{u}^i$ or $U^{-i}(p_M, p_M) \geq \underline{u}^{-i}$. Suppose for example the latter, then $v_{**}^i \geq V^i(p_M, p_M) \geq \pi(p_M, p_M)$.

Case $q_* = p_M$ and $\dot{p} \neq p_M$

1. We need to consider only $U^i(p_M, p_M) < \underline{u}^i$. Otherwise, $v_{**}^{-i} \geq V^{-i}(p_M, p_M) \geq \pi(p_M, p_M)$.

2. Then, by Lemma A2: $U^{-i}(p_M, p_M) \geq \underline{u}^{-i}$ and, thus, $\underline{u}^i > U^i(p_M, p_M) \geq \pi(p_M, p_M)$.

3. In addition, $p_* \neq p_M$. Otherwise, consider algorithm $s = \{x \to x\}$. For this algorithm $v^{-i}(s) \geq \pi(p_M, p_M)$ and, thus, $\underline{u}^i \leq u^i(s) \leq \pi(p_M, p_M)$.

4. Let us show that $U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^{-i}$

   - If not, then by Lemma A1: $U^i(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^i$ and, thus:

$$U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \quad \begin{aligned} &\geq (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot v^{-i}(p_{M-\Delta}, p_{M-\Delta}) \\ &\geq \pi(p_{M-\Delta}, p_{M-\Delta}) \end{aligned} \quad (77)$$

   - However, by Lemma A3: $\underline{u}^{-i} \leq U^{-i}(q_*, p_*) = \pi(p_M, p_*) < \pi(p_{M-\Delta}, p_*) \leq \pi(p_{M-\Delta}, p_{M-\Delta})$.

5. Finally since $U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^{-i}$:

$$v_{**}^i \geq V^i(p_{M-\Delta}, p_{M-\Delta}) \quad \begin{aligned} &\geq (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \underline{u}^i \\ &\geq (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \pi(p_M, p_M) \end{aligned} \quad (78)$$

Case $q_* = p_M$ and $\dot{p} = p_M$

1. In this case the monopoly price pair (p_M, p_M) is an absorbing state.

   - By Lemma A3: $\underline{u}^{-i} \leq U^{-i}(q_*, p_*) \leq \pi(p_M, p_M)$. Similarly, $\underline{u}^i \leq \pi(p_M, p_M)$.
   - By Lemma A2 for one of the sellers (e.g. $-i$): $U^{-i}(p_M, p_M) \geq \underline{u}^{-i}$, which implies: $U^i(p_M, p_M) = (1-\beta) \cdot \pi(p_M, p_M) + \beta \cdot v^i(p_M, p_M) \geq \pi(p_M, p_M) \geq \underline{u}^i$.
   - Finally, $v^{i,-i}(p_M, p_M) \geq \pi(p_M, p_M)$ and $u^{i,-i}(p_M, p_M) \geq U^{i,-i}(p_M, p_M) \geq \pi(p_M, p_M)$, which implies after (p_M, p_M) are subsequent price pairs are also (p_M, p_M).

2. Thus, $v^{i,-i}(p_M, p_M) = u^{i,-i}(p_M, p_M) = V^{i,-i}(p_M, p_M) = U^{i,-i}(p_M, p_M) = \pi(p_M, p_M)$.

3. Suppose $U^i(p_{M-\Delta}, p_{M-\Delta}) = (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot v^i(p_{M-\Delta}, p_{M-\Delta}) > \pi(p_M, p_M)$

   - Since $v^i(p_{M-\Delta}, p_{M-\Delta}) \leq V^i(p_*, q_*)$ then $p_* \neq p_M$.
   - If $U^{-i}(p_{M-\Delta}, p_{M-\Delta}) < \underline{u}^{-i}$ then by lemma 1: $U^i(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^i$.
     - As a result, $U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \geq \pi(p_{M-\Delta}, p_{M-\Delta})$.
     - However, by Lemma A3: $\underline{u}^{-i} \leq \pi(q_*, p_*) < \pi(p_{M-\Delta}, p_{M-\Delta})$.

35

- Finally, if $U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^{-i}$:

$$
\begin{aligned}
V^i(p_{M-\Delta}, p_{M-\Delta}) \; &= (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + u^i(p_{M-\Delta}, p_{M-\Delta}) \\
&\geq (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + U^i(p_{M-\Delta}, p_{M-\Delta}) \qquad (79) \\
&\geq (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \pi(p_M, p_M)
\end{aligned}
$$

4. Thus, we can assume that $U^{i,-i}(p_M, p_M) \geq U^{i,-i}(p_{M-\Delta}, p_{M-\Delta})$, which also implies that $v^{i,-i}(p_{M-\Delta}, p_{M-\Delta}) \geq \pi(p_M, p_M)$.

5. Consider $v^i(p_{M-\Delta}, p_{M-\Delta}) = V^i(p', q')$:

- If $q' \neq p_M$ then $v^i_{**} \geq V^i(p', q') \geq \pi(p_M, p_M)$.
- If $q' = p_M$ and $p' < p_M$ then

$$
\begin{aligned}
V^i(p', q') &= (1-\beta) \cdot \pi(p', p_M) + \beta \cdot u^i(p', p_M) \\
U^{-i}(q', p') &= (1-\beta) \cdot \pi(p_M, p') + \beta \cdot v^{-i}(p_M, p')
\end{aligned}
\qquad (80)
$$

  - Since $U^{-i}(q', p') \geq U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \geq (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \pi(p_M, p_M)$ and $\pi(p_{M-\Delta}, p_{M-\Delta}) > \pi(p_M, p')$ then $v^{-i}(p_M, p') > \pi(p_M, p_M)$.
  - However, by Lemma 2: $u^i(p', p_M) \geq V^i(p', q') \geq \pi(p_M, p_M)$, which implies $v^{-i}(p_M, p') \leq \pi(p_M, p_M)$.

- Finally, if $v^{i,-i}(p_M, p_M) = V^{i,-i}(p_M, p_M) = \pi(p_M, p_M)$ then

$$
V^{i,-i}(p_{M-\Delta}, p_{M-\Delta}) = (1-\beta) \cdot \pi(p_{M-\Delta}, p_{M-\Delta}) + \beta \cdot \pi(p_M, p_M) \qquad (81)
$$

and by Lemma A1 either $U^i(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^i$ or $U^{-i}(p_{M-\Delta}, p_{M-\Delta}) \geq \underline{u}^{-i}$.

Finally to complete the proof of Lemma 4 and hence Theorem 2, we present the proofs of Lemmas A1, A2 and A3.

Proof of Lemma A1

Suppose $U^{-i}(q, p) < \underline{u}^{-i}$ and $q \neq q_*$. If we consider algorithm $s^i = \{s^i_* | q \to p\}$ for seller i then $v^i(s^i) = v^i_*$, which implies that:

$$
U^i(p, q) = (1-\beta) \cdot \pi(p, q) + \beta \cdot v^i(p, q) \geq (1-\beta) \cdot \pi(p_C, p_C) + \beta \cdot v^i_* \geq \underline{u}^i. \qquad (82)
$$

Proof of Lemma A2

1. By Lemma A1 we only need to consider the case in which $q_* = \dot{p} = p_M$.

2. Assume $U^i(p_M, p_M) < \underline{u}^i$ and $U^{-i}(p_M, p_M) < \underline{u}^{-i}$ then in any equilibrium sequence we should never observe a monopoly price pair (p_M, p_M).

36

3. Consider some arbitrary equilibrium sequence $s_0, s_1, s_2, \ldots$ Suppose "odd" algorithms correspond, for example, to seller i and, in addition:

$$
\begin{aligned}
\xi(s_0, s_1) &= (x, p_M) & x \neq p_M \\
\xi(s_1, s_2) &= (y, z) & z \neq p_M
\end{aligned}
\tag{83}
$$

then the monopoly price pair (p_M, p_M) is possible.

- To prove this first note that:

$$
\underline{u}^i \leq u^i(s_0) = (1 - \beta) \cdot \pi(p_M, x) + \beta \cdot V^i(y, z) < (1 - \beta) \cdot \pi(p_M, p_M) + \beta \cdot v^i_{**} \tag{84}
$$

- Now consider the following two unrelated algorithms:

$$
\begin{aligned}
s^{-i} &= \{\underline{s}^{-i} | p_M \to p_M\} \\
s^i &= \{s^i_{**} | p_M \to p_M\}
\end{aligned}
\tag{85}
$$

- By responding with s^i to s^{-i} the seller i would get at least $(1 - \beta) \cdot \pi(p_M, p_M) + \beta \cdot v^i_{**} > \underline{u}^i$ while if the next price pair is not (p_M, p_M) they would get at most $\underline{u}^i$.

4. Consider an equilibrium sequence $s_0, s_1, s_2, \ldots$ starting with $s_0 = s_T = \{x \to x \ for \ any \ x\}$. In this sequence:

- $\xi(s_0, s_1) = (x_0, x_0) \neq (p_M, p_M)$.
- $\xi(s_1, s_2) = (x_1, p_M)$: otherwise we can replace s_1 with $s'_1 = \{s_1 | p_M \to p_M\}$ and improve payoffs of both sellers in the previous period.
- $\xi(s_k, s_{k+1}) = (x_k, p_M)$ for $k > 1$: otherwise by previous conjecture there is should exist a monopoly price pair.
- But, then this sequence cannot be an equilibrium. The seller with the highest $\pi(x_k, p_M)$ would find optimal to keep the algorithm s_k.

Proof of Lemma A3

1. Consider $V^i(p_*, q_*)$ and the sequence of price pairs that follows (p_*, q_*) in equilibrium:

$$
(p_*, q_*) \to (p', q') \to (p'', q'') \to \ldots
\tag{86}
$$

2. By Lemma 2: $V^i(p_*, q_*) \leq U^i(p', q')$ and $V^{-i}(q', p') \leq U^{-i}(q'', p'')$.

3. In addition for $q_* = p_M$: $U^{-i}(q_*, p_*) \leq V^{-i}(q', p')$

- $U^i(p_*, q_*) = (1 - \beta) \cdot \pi(p_*, q_*) + \beta \cdot v^i_* \geq \underline{u}^i$ since $\underline{u}^i \leq u^i(\{x \to p_c\}) \leq (1 - \beta) \cdot \pi(p_C, p_C) + \beta \cdot v^i_*$.
- As a result, $V^{-i}(q', p') \geq V^{-i}(q_*, p_*) = (1 - \beta) \cdot \pi(q_*, p_*) + \beta \cdot u^{-i}(q_*, p_*) \geq (1 - \beta) \cdot \pi(q_*, p_*) + \beta \cdot U^{-i}(q_*, p_*)$.

37

- Finally, since $U^{-i}(q_*, p_*) = (1 - \beta) \cdot \pi(q_*, p_*) + \beta \cdot V^{-i}(q', p')$ then $V^{-i}(q', p') \geq \pi(q_*, p_*)$ and, thus, $V^{-i}(q', p') \geq U^{-i}(q_*, p_*)$.

4. If $U^i(p', q') \leq V^i(p'', q'')$ then since (p_*, q_*) is the first-best price pair for i:

$$
\begin{aligned}
V(p_*, q_*) &= U(p', q',) = V(p'', q'') = \pi(p_*, q_*) = \pi(p', q') \\
U(q_*, p_*) &= U(q', p') = U(q'', p'') = \pi(q_*, p_*) = \pi(q', p')
\end{aligned}
\tag{87}
$$

- Suppose, there exists $\tilde{p}$ such that $\pi(\tilde{p}, q_*) \geq \pi(p_*, q_*)$ and $\pi(q_*, \tilde{p}) > \pi(q_*, p_*)$.
- Consider $U^{-i}(q_*, \tilde{p}) = (1 - \beta) \cdot \pi(q_*, \tilde{p}) + \beta \cdot v^{-i}(q_*, \tilde{p})$. Since $-i$ could use the following algorithm:

$$
s^{-i} = \{p_* \to q_*, \tilde{p} \to q_*, x \to \underline{s}^{-i}(x) \, for \, x \neq p_*, \tilde{p}\}
\tag{88}
$$

then $v^{-i}(q_*, \tilde{p}) \geq v^{-i}(s^{-i}) \geq \pi(p_*, q_*) \geq \underline{u}^{-i}$ and, thus, $U^{-i}(q_*, \tilde{p}) \geq \underline{u}^{-i}$. But, then (p_*, q_*) is not the first best price pair.

5. If $U^i(p', q') > V^i(p'', q'')$ the by lemma 2 either $(p'', q'') = (p_*, q_*)$ or $(p'', q'') = (p_{**}, q_{**})$ and $q' = q_*$.

6. If $(p'', q'') = (p_*, q_*)$ then

$$
\begin{aligned}
U^{-i}(q_*, p_*) &= U^{-i}(q', p') = \pi(q_*, p_*) = \pi(q', p') \\
\pi(p', q') &> \pi(p_*, q_*)
\end{aligned}
\tag{89}
$$

However, then $V^i(p', q') > V^i(p_*, q_*)$ and $U^{-i}(q', p') \geq U^{-i}(q_*, p_*) \geq \underline{u}^{-i}$ and, thus (p_*, q_*) is not the first best price pair.

7. If $(p'', q'') = (p_{**}, q'_{**})$ and $q' = q_*$ then:

- $U^i(p', q') > V^i(p'', q'')$ implies that $\pi(p_*, q_*) < \pi(p', q_*)$.
- If $\pi(q_*, p') > \pi(q_*, p_*)$ then

$$
\begin{aligned}
U^{-i}(q_*, p') &\geq (1 - \beta) \cdot \pi(q_*, p') + \beta \cdot V(q_*, p') > U^{-i}(q_*, p_*) \geq \underline{u}^{-i} \\
V^i(p', q_*) &\geq (1 - \beta) \cdot \pi(q_*, p') + \beta \cdot U(p', q_*) > V^i(p_*, q_*)
\end{aligned}
\tag{90}
$$

and, thus, again (p_*, q_*) is not the first best price pair.

8. Finally, let us prove that $V^i(p_*, q_*) = \pi(p_*, q_*)$ and $U^{-i}(q_*, p_*) = \pi(q_*, p_*)$.

- Since $U^i(p_*, q_*) \geq \underline{u}^i$ then $V^{-i}(q', p') \geq V^{-i}(q_*, p_*) \geq \pi(q_*, p_*)$ which in turns implies that $U^{-i}(q_*, p_*) \geq \pi(q_*, p_*)$. In addition, $V^i(p_*, q_*) \geq \pi(p_*, q_*)$.
- By assumption 0–(iv) from section 2.1: $(\pi(p_*, q_*), \pi(q_*, p_*))$ lies on the convex Hull of the payoffs set[18], which means the above inequalities can only hold if all subsequent price pairs are (p_*, q_*).

---

[18]Even if $p_* < \arg\max \tilde{p}\pi(\tilde{p}, p_M)$ then (p_*, q_*) is the point with the highest possible payoff on one seller.

Proof of Lemma 5

1. Similarly, as in the proof of Proposition 3 we can show it is never optimal to respond with an algorithm that would generate an immediate payoff of $\pi(p_M, p_C)$ and that $f^i(s_T) \neq s_C$.

2. Consider some equilibrium sequence: $s_0, s_1, s_2, s_3, \ldots$

3. The following two equilibrium sequences are not possible:

   - $s_T, s_R, s_C, s_3, \ldots$: "odd" seller can deviate to $s_1' = s_T$ and $s_3' = s_3$ (note that s_3 is either s_C or s_T) and improve their payoff

   - $s_R, s_R, s_C, s_3, \ldots$: "odd" seller can deviate to $s_1' = s_C$ and $s_3' = s_3$

4. Suppose $f^i(s_T) = s_R$ then given the above only the following equilibrium sequences starting with s_T chosen by seller $-i$:

   - $s_T, s_R, s_T, s_R, \ldots$: "odd" seller can deviate to $s_1' = s_T$ and $s_3' = s_3 = s_R$

   - $s_T, s_R, s_R, s_R, \ldots$: "odd" seller can deviate to $s_1' = s_T$ and then keep this algorithm

   - $s_T, s_R, s_R, s_T, s_R, s_T, s_R, \ldots$: same as above starting from s_3

   - $s_T, s_R, s_R, s_T, s_T, \ldots$ or $s_T, s_R, s_R, s_T, s_M, \ldots$: "odd" seller can deviate to $s_1' = s_T$

Proof of Proposition 2

1. Because of Lemma 5 the first five items in the proof of proposition 1 are still valid.

2. Consider one seller and suppose they are responding to algorithm s_R chosen by the competitor.

   - If the seller chooses s_C then the expected payoff is :

$$(1-\beta)\cdot\pi(s_C, s_R) + \beta(1-\beta)\cdot\pi(s_C, s_T) + \beta^2\cdot\pi(p_M, p_M) = (1-\beta)\cdot(1+x)+\beta^2 \quad (91)$$

   - If the seller chooses s_T then the expected payoff is:

$$(1-\beta)\cdot\pi(s_T, s_R) + \beta\cdot\pi(p_M, p_M) \quad (92)$$

   which is equal to (for two options of resolving cycles):

$$\begin{aligned} (minimum\,price) &= \beta \\ (average\,payoff) &= (1-\beta)\cdot\left(\tfrac{1}{2}+\tfrac{x-y}{4}\right)+\beta \end{aligned} \quad (93)$$

   - Responding with s_T could be preferable to s_C only for "average payoff" case and only if

$$(1-\beta)\cdot(1+x)+\beta^2 \leq (1-\beta)\cdot\left(\frac{1}{2}+\frac{x-y}{4}\right)+\beta \qquad \Leftrightarrow \quad (94)$$

39

$$y \leq 4\beta - 2 - 3x \tag{95}$$

which is only possible if $\beta \geq \frac{1}{2}$ and $x < \frac{4\beta-2}{3} < \beta$ (thus, only type I equilibrium could be affected).

Part B: Non-existent "second best" price pair

In this section we consider a special case, in which the "second best" price pair from Definition 5 does not exist for at least one of the sellers. We can show that in this case for an arbitrary equilibrium sequence $s_0, s_1, s_2, \ldots$ the continuation payoffs for at least one of the seller will exceed the monopoly payoff $\pi(p_M, p_M)$. The steps bellow prove this result. For the "first best" pair pairs we use the same notation equation (75) as in the proof of Lemma 4.

1. Suppose, for example, $U^{-i}(q, p) < \underline{u}^{-i}$ for any $p, q \neq q_*$. This implies $U^{-i}(q_*, p) \geq \underline{u}^{-i}$ for any p (seller i can construct an algorithm which would lead to prices (q_*, p)).

2. As a result, Lemma 2 for seller i means either $v^i(s_{k+1}) \geq u^i(s_k)$ or $v_i(s_{k+1}) = v_*^i$. Thus, the continuation payoffs for seller i in the any equilibrium sequence either become constant (the we can use Lemma 3) or reaches v_i^*.

3. If $q_* = p_M$ then $v_*^i \geq V^i(p_M, p_M) \geq \pi(p_M, p_M)$. Thus, assume $q_* \neq p_M$.

4. If $\dot{p} \neq p_M$ then $v_{**}^{-i} \geq V^{-i}(p_M, p_M) \geq \pi(p_M, p_M)$. Thus, assume $\dot{p} = p_M$. (Since $U^i(p_M, p_M) \geq \pi(p_M, p_M)$, the "second best" price pair necessarily exists for seller $-i$).

5. Since $U^{-i}(\dot{q}, \dot{p}) = (1 - \beta) \cdot \pi(\dot{q}, \dot{p}) + \beta \cdot v_*^{-i} \geq u^{-i}(\{x \to p_C\}) \geq \underline{u}^{-i}$ then $\dot{q} = q_*$.

6. Finally for any p: $U^{-i}(q_*, p) = (1 - \beta) \cdot \pi(q_*, p) + \beta \cdot v^{-i}(q_*, p)$. Then, either $p = \dot{p}$ and $v^{-i}(q_*, p) = v_*^{-i}$ or $p \neq \dot{p}$. This together with Lemma 2 implies that in any equilibrium sequence either continuation payoffs for seller $-i$ converge to a constant or they become above $v_*^{-i} \geq V^{-i}(p_M, p_M) \geq \pi(p_M, p_M)$.

Part C: Explicitly including pre-convergent prices into the calculations

Here we define the precise model and calculate expected payoffs within this model, and study the Markov equilibria in this updated model. A key difference is that the payoff relevant variable now is not only the algorithm of the opponent, but also the next price that this algorithm chooses. Then, we formally establish and prove the claim made in Section 5.2. For simplicity and to relax assumptions piecemeal, we continue to assume that the seller with the revision opportunity deciphers the opponent's algorithm at a very small cost, see part E of the Appendix for a justification.

40

Precise Model

Similar to our main model we assume that prices change in the discrete moments $0, 1 \cdot \Delta t, 2 \cdot \Delta t, \ldots$ with subsequent prices determined by

$$p^A_{t+\Delta t} = s^A(p^B_t)$$
$$p^B_{t+\Delta t} = s^B(p^A_t) \tag{96}$$

where s^A and s^B are the current algorithms of the two sellers. Customers arrive continuously and randomly (also potentially between price adjustments) with Poisson intensity λ. The continuously compounded discount rate is r. Before each price adjustment the seller who was not the last to adjust their algorithm gets the opportunity to adjust their algorithms with probability $(1 - e^{-\mu \Delta t})$. When this happens this seller can also choose prices for the current and for the next period.

As before, we *index* all moments in time when one of the sellers adjusting their algorithm with $k = 0, 1, 2, \ldots$. and consider some arbitrary adjustment moment k that happens at time t_k. We express the "precise" expected payoff of the seller adjusting their algorithm, $\tilde{U}_k$, and of the other seller, $\tilde{V}_k$, as follows (note that $\int_0^{\Delta t} \lambda \pi_{k,i} e^{-rt} dt = \frac{\lambda}{r}(1 - e^{-r\Delta t})\pi_{k,i}$):

$$\tilde{U}_k = \sum_{i=0}^{\infty} e^{-(r+\mu)i\Delta t} \cdot \left[\frac{\lambda}{r}(1 - e^{-r\Delta t})\pi_{k,i} + (1 - e^{-\mu\Delta t})\tilde{V}_{k+1} \right]$$

$$\tilde{V}_k = \sum_{i=0}^{\infty} e^{-(r+\mu)i\Delta t} \cdot \left[\frac{\lambda}{r}(1 - e^{-r\Delta t})\bar{\pi}_{k,i} + (1 - e^{-\mu\Delta t})\tilde{U}_{k+1} \right] \tag{97}$$

where $\pi_{k,i}$ and $\bar{\pi}_{k,i}$ are the expected payoff from one customer for the corresponding seller in the interval $(t_{k+i\cdot\Delta t}, t_{k+(i+1)\cdot\Delta t})$, in which the prices are constant.

If we denote average weighted per period payoffs of each sellers as:

$$\pi_k = \frac{\sum_{i=0}^{\infty} e^{-(r+\mu)i\Delta t}\pi_{k,i}}{\sum_{i=0}^{\infty} e^{-(r+\mu)i\Delta t}} \qquad \bar{\pi}_k = \frac{\sum_{i=0}^{\infty} e^{-(r+\mu)i\Delta t}\bar{\pi}_{k,i}}{\sum_{i=0}^{\infty} e^{-(r+\mu)i\Delta t}} \tag{98}$$

and normalize payoffs as $\tilde{u}_k = \frac{\tilde{U}_k}{\lambda/r} \cdot e^{-\mu\Delta t}$ and $\tilde{v}_k = \frac{\tilde{V}_k}{\lambda/r} \cdot e^{-\mu\Delta t}$ then we get the following equations describing the "precise" normalized payoffs:

$$\tilde{u}_k = (1 - \tilde{\beta}) \cdot \pi_k + \tilde{\beta} \cdot \tilde{v}_{k+1}$$
$$\tilde{v}_k = (1 - \tilde{\beta}) \cdot \bar{\pi}_k + \tilde{\beta} \cdot \tilde{u}_{k+1} \tag{99}$$

Here $\tilde{\beta} = \frac{1-e^{-\mu\Delta t}}{1-e^{-(\mu+r)\Delta t}}$ is the effective discount factor.

Equations (99) are similar to equations (12) describing the evolution of normalized payoffs in the main model. In addition, $\tilde{\beta} \to \beta = \frac{\mu}{\mu+r}$ as $\Delta t \to 0$. Finally, if $\pi(s_k, s_{k-1})$ and $\bar{\pi}(s_{k-1}, s_k)$ are the limiting expected payoffs from one customer for the two algorithms s_k and

s_{k-1}[19], then it follows directly from equations ([98]) that:

$$|\pi_k - \pi(s_k, s_{k-1})| \leq \Delta\pi_{max} \cdot \left(1 - e^{-(r+\mu)\cdot N_k \Delta t}\right)$$
$$|\bar{\pi}_k - \bar{\pi}(s_{k-1}, s_k)| \leq \Delta\pi_{max} \cdot \left(1 - e^{-(r+\mu)\cdot N_K \Delta t}\right) \tag{100}$$

where $\Delta\pi_{max} = max_{p,q,p',q'} |\pi(p,q) - \pi(p',q')|$ and N_k is the length of the converge period for the current pair of algorithms (which cannot be higher than the number of possible prices squared). Thus, for the same behavior of the sellers the expected payoffs in the "precise" model ([99]) would converge to the expected payoffs of the "approximate" limiting model ([12]) as $\Delta t \to 0$.[20]

Equilibria in the Precise Model

For sufficiently small Δt the initial price would not have any significant effect on the expected payoff of the seller adjusting their algorithm—by choosing the appropriate initial prices they can always get the continuation sequence with the highest expected payoffs. However, it could be used as a coordination device allowing the seller to select one of the possible best responses with practically identical payoffs. In order not to do deal with such situations we assume below that all payoffs in the payoff matrix $\pi(p,q)$ are distinct, that is:

$$\pi(p,q) = \pi(p',q') \quad \Leftrightarrow \quad (p,q) = (p',q') \tag{101}$$

Now consider some arbitrary Markov best response functions for two sellers:

$$\{f^i(s,p), \ p_0^i(s,p), \ p_1^i(s,p)\} \qquad i = A, B \tag{102}$$

Here $f^i(s,p)$ is the algorithm with which seller i would respond and $p_0^i(s,p)$ and $p_1^i(s,p)$ are the initial two prices chosen. Note, the state variable here includes, in addition to the current algorithm of the opponent s, the first price p that the opponent's algorithms will choose. For the pair of responses ([102]) we can calculate the expected payoffs of the seller adjusting their algorithm $\tilde{u}^i(s,p)$ and of the other seller $\tilde{v}^i(s,p)$ for given values of s and p. Also, denote by $u^i(s,p)$ and $v^i(s,p)$ the limits of these expected payoffs as $\Delta t \to 0$ (for the same response functions of the sellers ([102])):

$$u^i(s,p) = \lim_{\Delta t \to 0} \tilde{u}(s,p)$$
$$v^i(s,p) = \lim_{\Delta t \to 0} \tilde{v}(s,p) \tag{103}$$

Proposition 3. *Suppose that for any $\varepsilon > 0$ there exists $\Delta t = \Delta t_\varepsilon < \varepsilon$ such that the pair of responses from ([102]) is an equilibrium in the "precise model" then for any β (except for possibly some finite set of values):*

---

[19]To be precise these payoffs also depend on the initial price sequence chosen by the seller adjusting their algorithms. We skip this here for more concise notation.

[20]We again assume sellers cannot respond with an algorithm that would generate a cycle. This assumption can be easily relaxed though. If we resolve cycles in the main model as equivalent to obtaining the average payoffs over the cycle, then similarly the "precise" payoffs π_k and $\bar{\pi}_k$ would converge to these averages as $\Delta t \to 0$.

1. $u^i(s, p) = u^i(s, p')$ and $v^i(s, p) = v^i(s, p')$ for any prices p and p'

2. For arbitrary $p^i(s)$ the strategies

$$\tilde{f}^i(s) = f^i(s, p^i(s)) \tag{104}$$

are an equilibria in the "main model" with expected payoffs $u^i(s) = u^i(s, p')$ and $v^i(s) = v^i(s, p')$ for any p'.

Proposition 3 shows that if some pair of best responses (102) survives as an equilibrium for an arbitrary small Δt then we can construct an equilibrium in the main model by resolving the possible dependence on price arbitrarily (equation (104)). Moreover, the expected payoffs in the limiting main model would be equal to the limits of the expected payoffs of the precise model as $\Delta t \to 0$. Since there is a finite number of possible response pairs (102) the next corollary directly follows.

Corollary 3. *For any $\beta = \frac{\mu}{\mu + r}$ (except for possibly some finite set) and for sufficiently small Δt if we take any Markov equilibrium in the "precise" model and resolve dependency on the initial price arbitrarily then the obtained strategies would be an equilibrium in the "approximate" limiting model.*

Thus, if we observe only collusive equilibria in the main model, then all equilibria in the "precise model" for sufficiently small Δt would also be collusive (except for possibly some non-generic subset of discount factors β).

Proof of Proposition 3

1. Consider one of the sellers choosing the best response for some s and p. For sufficiently small Δt the seller would always respond with an algorithm s' and initial price sequence that would generate the highest possible limiting expected payoff:

$$u(s, p) = \max_{p', s'(.)} (1 - \beta) \cdot \pi(p', s(p')) + \beta \cdot v(s', p') \qquad s.t. \ s'(s(p')) = p' \tag{105}$$

which does not depend on p.

2. If if there are several different p' and s' that generate the highest limiting payoff for the seller adjusting their algorithm then with assumption (101) except for possibly finite set of discount factors β these responses also generate exactly the same limiting expected payoff for the other seller:

$$v(s, p) = (1 - \beta) \cdot \pi(s(p'), p') + \beta \cdot u(s', p') \tag{106}$$

3. As in the main model, for sufficiently small Δt the seller adjusting the algorithm would also choose the initial prices that maximizes the current expected payoff (since any future

43

responses of the opponent generate the same limiting expected payoff irrespective of the future initial prices).

4. Finally, the strategies (104) are an equilibrium in the "approximate" limiting model since any profitable deviation would also be profitable in the "precise" model for sufficiently small Δt.

Part D: _Example with unequal equilibrium_

In this section we construct an example of an unequal equilibrium, in which prices can converge to a collusive that is not (p_M, p_M). For this example we will use the discrete choice demand function with the following expected payoff from one customer:

$$\pi(p, q) = p \cdot \frac{e^{-bp}}{a + e^{-bp} + e^{-bq}} \tag{107}$$

We choose parameters $a \approx 0.0158$ and $b \approx 0.4760$ so that the monopoly price is $p_M = 8$ and the competitive price is $p_C = 4$. Each seller can chose any prices from the set $\{4, 5, 6, 7, 8\}$. The payoff matrix for these prices with the above parameters is shown in Table 6.

To describe an equilibrium instead of specifying all algorithms that are used (which would be quite complicated) we instead specify the transition matrices, which show which price pair could follow another price pair. Later we can verify that the chosen transition matrices can indeed be sustained in an equilibrium. More specifically, consider the best algorithm that could follow prices (p, q) for seller i (from equation (19)):

$$s^i(p, q) = \arg \max_{s: q \to p} v^i(s) \tag{108}$$

Now denote by $\phi^i(p, q)$ the price pair that would follow (p, q) on the equilibrium sequence starting with $s^i(p, q)$ or more precisely

$$\phi^i(p, q) = \xi^{-i}(s^i(p, q), f^{-i}(s^i(p, q))) \tag{109}$$

If we know the transition matrices $\phi^i(p, q)$ for each seller we can completely describe how prices will evolve and we can calculate the expected payoffs matrices $v^i(p, q)$, $u^i(p, q)$, $V^i(p, q)$, $U^i(p, q)$. Finally, using Proposition 1 we verify that the specified transitions are feasible and optimal, and, thus, they indeed correspond to an equilibrium. The specific algorithms necessary to sustain such equilibrium can then be recovered from equation (31).

Consider the transition matrices specified in Table 7. We verified that the specified transition matrices can be sustained in an equilibrium for $\beta = 0.9, 0.95, 0.999$ and, also, for any β sufficiently close 1.[21] For the specified transition matrices, depending on the initial condition

---

[21]We performed the calculations numerically with a Python script. Nevertheless, when comparing two expected payoffs we required for the difference to be higher than the specified threshold, which was much higher than possible mistakes in the numerical calculations. Thus, we believe the proof is strict. To show

$A\backslash B$	8	7	6	5	4
8	2.95, 2.95	2.41, 3.39	1.86, 3.61	1.36, 3.54	0.95, 3.19
7	3.39, 2.41	2.87, 2.87	2.29, 3.16	1.74, 3.21	1.25, 2.97
6	3.61, 1.86	3.16, 2.29	2.64, 2.64	2.08, 2.79	1.55, 2.68
5	3.54, 1.36	3.21, 1.74	2.79, 2.08	2.30, 2.30	1.80, 2.32
4	3.19, 0.95	2.97, 1.25	2.68, 1.55	2.32, 1.80	1.90, 1.90

Table 6: Expected payoffs from one customer in an unequal equilibrium.

Seller A

$A\backslash B$	8	7	6	5	4
8	8, 8	8, 7	8, 6	8, 6	8, 7
7	7, 8	8, 8	8, 7	8, 7	8, 7
6	5, 4	7, 8	8, 8	8, 8	8, 8
5	5, 4	7, 8	7, 8	8, 8	7, 8
4	5, 4	7, 8	7, 8	7, 8	8, 8

seller B

$A\backslash B$	8	7	6	5	4
8	8, 8	7, 7	7, 7	7, 7	7, 7
7	7, 8	8, 8	8, 8	8, 8	8, 8
6	7, 8	7, 8	8, 8	8, 8	8, 8
5	7, 8	7, 8	7, 8	8, 8	7, 8
4	7, 8	7, 8	7, 8	7, 8	8, 8

Table 7: Price transition matrices (the first price in each pair corresponds to seller A)

we end up either with monopoly prices $(8,8)$ or with an unequal outcome $(7,8)$ in which seller A gets a payoff above the monopoly. In order to sustain prices $(7,8)$ in an equilibrium seller A can use the following algorithm:

$$s^A = \{8 \to 7, \ 7 \to 6, \ 6 \to 5, \ 5 \to 4, \ 4 \to 5\} \tag{110}$$

In response to such algorithm seller B cannot get more that $\pi(8,7)$ immediately and if seller B believes that seller A will keep this algorithm they do not have any incentives to continue with any other price pair. While for A it is optimal to keep this algorithm because it guarantees higher than the monopoly expected payoff.

Part E: _Modeling experimentation_

In this section we discuss in more detail how sellers learn the opponent's algorithm. In particular, we argue that within our model, experimentation is relatively inexpensive (and converges to 0 as $\Delta t \to 0$), and, thus, the sellers will always choose to discover the opponent's algorithm (at least enough to make a decision). As a result, for Δt small enough, we can assume the seller to be best responding the opponent's algorithm.

Whenever a seller gets an opportunity to revise their algorithm, they can also spend a short period of time experimenting with the opponent's algorithm (by manually choosing an initial short sequence of prices) before committing to a new algorithm themselves.

Suppose $U^{observable}$ is the highest payoff the seller currently adjusting their algorithm could get if the opponent's algorithm is directly observable. Now, the considered seller can

---

that the transition matrices could be sustained as an equilibrium for β sufficiently close to 1 we created a slightly modified version of the script, which compared not the discounted expected payoffs but the sequences of payoffs.

experiment for the first K periods (for K different price levels) to fully recover the opponent's algorithm. Then, their ultimate payoff is then at least:

$$
\begin{aligned}
U &\geq \sum_{j=0}^{K-1} e^{-r \cdot j \Delta t} \cdot \frac{\lambda}{r} (1 - e^{-r\Delta t}) \cdot \pi(\dots, \dots) + e^{-r \cdot K \Delta t} \cdot U^{observable} \\
&\geq U^{observable} - \frac{\lambda}{r} \left(1 - e^{-r \cdot K \Delta t}\right) \Delta \pi_{max}
\end{aligned}
\tag{111}
$$

where $\Delta \pi_{max} = max_{p,q,p',q'} |\pi(p,q) - \pi(p',q')|$.[22] The second term in the above expressions is converging to 0 as $\Delta t \to 0$, thus, the actual expected payoffs approximate $U^{observable}$ for small Δt.[23] While if the seller does not experiment and chooses the response immediately, their expected payoff cannot be higher that $U^{observable}$ and could be sufficiently lower (even for small Δt) because the chosen response might not be optimal for the specific opponent's algorithm.

We should point out though that for general algorithms (with many possible internal states) experimentation could certainly be more costly and it might take more time. It might not even possible to fully decipher the opponent's algorithm and you might need to choose your algorithm with some remaining uncertainly. In addition, the opponent's algorithm might have some irreversible state transitions (e.g. some sort of grim-trigger algorithm), thus, it might not be obvious whether you want to experiment at all. It would, of course, be interesting to study experimentation with these general algorithms, but, also quite challenging. It is certainly beyond the scope of this paper, and we leave it for future research.

---

[22]In the expression we also ignore the probability simultaneous algorithm adjustments, which for small Δt is a small probability event, and, thus, similarly would not materially affect the ultimate payoff.

[23]Since the experimentation phase is not longer than K prices adjustment cycles, it could be even shorter the period of convergence to a fixed price pair, which we already ignore in our model by Assumption 1.

References

D. Abreu and A. Rubinstein. The structure of Nash Equilibrium in repeated games with finite automata. *Econometrica*, 56(6):1259–1281, 1988.

D. Abreu, D. Pearce, and E. Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, 1990.

J. Asker. A study of the internal organization of a bidding cartel. *American Economic Review*, 100(3):724–762, 2010.

J. Asker, C. Fershtman, and A. Pakes. Artificial intelligence and pricing: The impact of algorithm design. UCLA, Tel Aviv and Harvard, 2021.

S. Assada, R. Clark, D. Ershovc, and L. Xud. Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. CESifo Working Paper No. 8521, 2021.

R. J. Aumann. Acceptable points in general cooperative n-person games. In A. W. Tucker and R. D. Luce, editors, *Contributions to the Theory of Games IV, Annals of Mathematics Study 40*. Princeton University Press, 1959.

Y. Awaya and V. Krishna. On communication and collusion. *American Economic Review*, 106(2):285–315, 2016.

M. Banchio and A. Skrzypacz. Artificial intelligence and auction design. Stanford University, 2022.

Z. Y. Brown and A. MacKay. Competition in pricing algorithms. *American Economic Journal: Microeconomics*, forthcoming, 2021.

E. Calvano, G. Calzolari, V. Denicolò, J. E. Harrington Jr., and S. Pastorello. Protecting consumers from collusive prices due to ai. *Science*, 370(6520):1040–1042, 2020a.

E. Calvano, G. Calzolari, V. Denicolò, and S. Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–3297, 2020b.

S. Chassang, K. Kawai, J. Nakabayashi, and J. Ortner. Robust screens for noncompetitive bidding in procurement auctions. *Econometrica*, 90(1):315–346, 2022.

J. W. Friedman. A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, 38(1):1–12, 1971.

E. J. Green and R. H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100, 1984.

K. T. Hansen, K. Misra, and M. M. Pai. Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Science*, 40(1):1–12, 2021.

J. E. Harrington Jr. Developing competition law for collusion by autonomous agents. *Journal of Competition Law Economics*, 14(3):331–363, 2018.

J. E. Harrington Jr. The effect of outsourcing pricing algorithms on market competition. Management Science, *forthcoming*, 2022.

J. Johnson, A. Rhodes, and M. Wildenbeest. Platform design when sellers use pricing algorithms. Cornell University and Toulouse School of Economics and Indiana University, 2021.

Y. Kamada and M. Kandori. Revision games. *Econometrica*, 88(4):1599–1630, 2020.

V. Kumar. Cartels in the kautiliya arthasastra. *Czech Economic Review*, 6(1):59–79, 2012.

R. Lagunoff and A. Matsui. Asynchronous choice in repeated coordination games. *Econometrica*, 65(6):1467–1477, 1997.

R. C. Marshall and L. M. Marx. *The Economics of Collusion: Cartels and Bidding Rings*. MIT Press, 2012.

E. Maskin and J. Tirole. A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica*, 56(3):549–569, 1988a.

E. Maskin and J. Tirole. A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, 56(3):571–599, 1988b.

D. A. Miller and J. Watson. A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350, 2013.

D. G. Pearce. Renegotiation-proof equilibria: Collective rationality and intertemporal cooperation. *Cowles Foundation Discussion Paper No. 855*, 1987.

M. Safronov and B. Strulovici. Contestable norms. University of Cambridge and Northwestern University, 2018.

B. Salcedo. Pricing algorithms and tacit collusion. University of Western Ontario, 2015.

T. Sugaya. Folk theorem in repeated games with private monitoring. Review of Economic Studies, *forthcoming*, 2020.