With the advent of the digital age, menu costs have been falling, yet little is known about how the ability to update prices more frequently affects competition between firms. While there is extensive theoretical literature on dynamic games, the Folk Theorem warns us that the predictive power of equilibrium reasoning is limited in dynamic environments with brief periods. Hence, empirical research into the issue of repricing is necessary.

This paper employs data on the pricing decisions made by third-party sellers on Amazon Marketplace to empirically assess and expand the theoretical predictions of the literature on dynamic pricing. We find that delegation of pricing to simple algorithms can facilitate tacit collusion by reducing the set of available strategies. Furthermore, the algorithms currently employed are more straightforward than the Markov-perfect strategies of Maskin and Tirole (1988) and emerge naturally as the result of a best-response process once the ability to reprice regularly is introduced into the marketplace. We build a model that suggests that if this best-response process were allowed to play out without intervention, the market would fully transition from static to dynamic pricing. Along the transition path, prices initially fall but eventually rise dramatically: the (average) equilibrium price between some repricing algorithms is the monopoly price.

This paper proceeds as follows. Firstly, we motivate our analysis. Secondly, we introduce the institutional setting. Thirdly, we distill our understanding of repricing using evidence from a regression discontinuity design and quasi-random variation in repricing strategy execution. The former establishes that turning on repricing lowers a merchant's prices without adversely affecting market prices, and the latter that resetting strategies (which regularly raise a merchant's prices) coax competitors to raise their prices. Fourthly, we develop a model of equilibrium in delegated strategies that reflects the limited set of strategies available to repricing merchants, simulate its implied long-run evolutionary dynamics and compare them to the data. Finally, we conclude.

# 1 Motivation

As the economy digitizes, it is becoming easier for firms to monitor rivals' prices and quickly respond to price changes. These trends are not limited to the world of e-commerce. Instead, they are general consequences of spreading information technology: retail chains are posting their prices online and brick-and-mortar stores are adopting electronic shelf labels. From 2013 to 2018, Tesco automatically matched its prices to competitors at the till (guardian.com).

As collusive schemes are easier to maintain when rivals' prices are observable and swift punishment for deviations possible, these developments are concerning (Rotemberg and Saloner, 1986; Green and Porter, 1984). Though better demand prediction need not facilitate collusion (Miklós-Thal and Tucker, 2019), the publication of firm-specific transaction prices in the Danish ready-mixed concrete market led to decreased price competition (Albaek et al., 1997). Furthermore, in a world with low search costs where offer-level elasticities can reach $-20$ (Ellison and Ellison, 2009), the payoffs to

An extended two-page abstract of this paper was previously published as Anonymous (2022).

collusion are high.

There is a recent worry that such collusion may be brought about by pricing algorithms. These algorithms, a newly burgeoning literature warns us, can *learn* to coordinate their actions (Salcedo, 2015; Calvano et al., 2020; Klein, 2021; Hansen et al., 2021) just like humans (Byrne and Roos, 2019) at least under some learning protocols (Asker et al., 2022).

Still, this warning is mostly speculative, and the concern is investigated using simulation studies or theoretical arguments. By contrast, this paper contributes an empirical study of tacitly collusive algorithms employed at-scale in practice. By necessity, then, we discuss incentives for the *adoption* of pre-existing algorithms but remain largely silent on how they were learnt. As pricing algorithms are themselves products, there may be pressure towards algorithms that successfully collude with their copies but are not easily exploited by non-algorithmic competitors (Harrington, forthcoming). Indeed, gasoline prices in Germany increase only when *both* members of a local duopoly adopt algorithmic pricing strategies (Assad et al., 2020). Similarly, the algorithms we find in this paper are based on undercutting strategies that make them hard to exploit.

Even disregarding collusion, algorithms' short-term commitment can yield price increases if they operate at asymmetric speeds (Brown and McKay, forthcoming). While our analysis abstracts from asymmetric speed, we also find that short-term commitment (to undercutting your opponent) and asymmetry (in available strategies) play key roles in pushing up prices. By contrast, managerial override – potentially a key feature of gasoline markets (Leisten, 2022) – is unlikely in our setting.

In our data, the initial effect of repricing was almost certainly pro-competitive as most merchants chose to reprice using simple undercutting strategies. More recently, however, the worry of facing other players employing similar undercutting strategies has led to repricing strategy innovation. Merchants are increasingly using tacitly collusive 'cycling' strategies that manifest in Edgeworth-like cycles: infrequent large price increases alternate with periods of frequent, stable, and small price decreases.

Edgeworth cycles emerge as Markov-Perfect equilibria in a dynamic price-setting model in which two firms of equal size and with constant unit costs set prices for perfect substitutes in alternating, discrete steps (Maskin and Tirole, 1988). The cycles are robust to alternative timing assumptions or differences in firm size (Eckert, 2003) and survive in a three-firm model, with imperfect substitutes or when marginal costs fluctuate (Noel, 2008). However, cycling is more likely in markets with high own-price elasticity and weak capacity constraints (Noel, 2008). Empirically, cycles have been found[1] in many gasoline markets (Noel, 2007; Wang, 2009a,b; Eckert, 2003) and online advertising auctions (Zhang, 2005; Edelman and Ostrovsky, 2007).

While reminiscent of Edgeworth cycles, our data are *not* generated by the MPE strategies proposed in Maskin and Tirole (1988). While MPE are a natural benchmark when humans are making pricing decisions, this ceases to be realistic when prices are changed as often as in our dataset. Instead, merchants delegate short-term pricing decisions to a computer (Chen et al., 2016). This requires expressing a strategy in the limited vocabulary of the available repricing interfaces,

[1] For an overview, see Noel (2018).

which makes a crucial difference to equilibrium outcomes. This delegation step is emphasized by Schlosser and Boissier (2017), who consider a merchants' (non-delegated) best-response to a specific (delegated) algorithm. Similarly, Popescu (2015) models how simple proportional repricing rules can emerge as the product of best response dynamics à la Milgrom and Roberts (1990). However, both stop short of modeling an equilibrium in delegated strategies, which is crucial to our finding that the average price over the cycle will be near monopoly.

## 2 Setting

With revenue of $470 billion in 2021, Amazon is one of the largest e-commerce platforms worldwide (amazon.com). Its success is built partly on allowing third-party sellers to list right alongside Amazon's offers: in 2017, more than half of units sold were from third-party sellers (amazon.com).

Amazon's web presence is organized around the concept of a (narrowly-defined) product: sellers do not create separate listings for the same product (as on, e.g., eBay). Instead, different offers are pooled on a unique product page. Amazon enforces this pooling by requiring all products to be listed under their UPC; listing a product without is not permitted.

As far as customers are concerned, the fact that there exist multiple offers for most products often goes ignored: about 83% of purchases go through the Buybox (repricerexpress.com), i.e., the framed section of the product page depicted in Figure 1a which prominently displays 'Buy Now' and 'Add to Cart' buttons. Amazon selects which merchant owns the Buybox using a (partially randomized) proprietary algorithm that loads heavily on which offer is the cheapest[2] (Lee and Musolff, 2021).

From a seller's perspective, the fact that offers are pooled on a product page means that winning the Buybox is crucial. Moreover, as sellers are aware that doing so is mostly a matter of having the lowest price, we might expect intense competition and prices close to marginal cost. Indeed, to a first approximation, we can think of the Buybox as emulating demand in a Nash-Bertrand game with homogeneous goods: it chooses the cheapest offer and exposes all demand to it. For tractability, below, we make precisely this assumption.

The Nash-Bertrand story suggests that markups on Amazon should be low throughout. However, in our data, the median markup[3] is 18.62%. Furthermore, there is considerable mass on even higher markups. So how do sellers avoid prices racing to the bottom?

The answer lies partially in the Marketplace Web Services (MWS) API, which allows sellers to reprice frequently at zero cost. Typically, a seller delegates repricing to an external repricing company. This company registers with MWS on the seller's behalf and subsequently is notified

[2] Intriguingly, such 'price directed prominence' and dynamic modifications have been found to curb collusion both theoretically and in simulation studies with Q-learning agents (Johnson et al., 2021).

[3] To be exact, the data contains the sales price $p$, the fees $f$ paid to Amazon for the fulfillment of the sale, and the seller-reported cost of the item $c$. We calculate the markup as $\mu = [p - (c + f)]/p$. The results are based on 3,279,322 sales observations with non-missing costs and fees. We do not observe and hence do not include in our calculation two other fees potentially paid by the seller: (i) the monthly subscription fee (which at $39.99 is negligible) and (ii) potential inventory storage fees (including monthly storage fees, long-term storage fees and FBA disposal order fees.)


(a) The Buybox (b) The Featured Merchants Section

Figure 1: How Offers Are Depicted on Product Page.

*Notes*: The left panel depicts the 'Buybox,' an area on the product page where the customer can directly purchase the product without making an explicit choice between different offers. The right panel illustrates that some additional offers may be listed directly on the product page.

by Amazon whenever there is some change to an offer on the seller's products. This notification is near-instant and contains all relevant details for the repricer: for the twenty most competitive offers (ranked by landed price, i.e., price plus shipping cost), it provides price, shipping cost, Buybox ownership, and more. The repricer takes this data, calculates a new price according to a predefined rule, and immediately sends it back to Amazon. We discuss the exact (and limited) rules available in more detail below.

Sellers pay repricers according to a typical 'digital' cost structure: they incur an average fixed menu cost of 0.04% of their revenue and pay no additional fee per repricing event. By comparison, supermarkets in 1992 spend 0.70% of revenue on menu costs (Levy et al., 1997). Thus, menu costs on Amazon differ from brick-and-mortar retail both in their structure (no marginal cost of repricing) and overall level (much lower).

Given the ability of sellers to reprice their products without incurring any (marginal) menu cost, economic theory suggests that the frequency of price changes should be high. The 'push' nature of the repricing company's pricing data allows us to confirm this hypothesis by observing the time elapsed between offer- and product-level price changes. The median offer updates on average every 25.71 hours. Furthermore, the median product's buybox price – arguably the most relevant price for consumers – updates on average every 10.62 hours. The median price spell duration in brick-and-mortar retail ranges from 1.5 months to 14.7 months (Bank, 2004, Table 3). In stark contrast, the median spell duration for the Amazon data is just 0.29 hours.[4]

While the median offer changes price frequently, the size of the change is typically small: the median (absolute) price change is $0.02. Furthermore, 78% of price changes are negative. These facts suggest that price changes are driven by competition: we expect to see frequent small negative

[4]Note that the concept of "median spell" pools price spells from different products.

price changes if firms are underbidding each other to gain the Buybox.

This paper will argue that undercutting is indeed common but accompanied by less frequent and more sizeable positive price changes. This pattern is illustrated in Figure 2, which depicts a typical Edgeworth-like cycle in our dataset. The product in question is a running shoe, and we focus on the two offers that obtain the Buybox at some point during our selected sample period. The figure shows the offer's price paths and current possession of the Buybox (in the color strip at the bottom of the graph). The bottom panel zooms in on the period marked by a grey backdrop in the top panel.

The general pattern of prices matches the predictions of Maskin and Tirole (1988): both prices decline over time until they jump back up again, one after the other. Furthermore, if we understand the Buybox as a proxy for demand, we see plenty of back-and-forths regarding which seller faces demand – again, as predicted.

However, once we examine the price paths in more detail in the bottom panel, some discrepancies with theory begin to emerge. Red has an inherent advantage: whenever the two sellers have the same price, Red is allocated the Buybox. On the other hand, Green must undercut Red to gain the Buybox. Furthermore, Red seems aware of this as he always chooses to match Green's price while Green is stuck strictly undercutting Red's price. Furthermore, at the bottom of the cycle, the war of attrition predicted by Maskin and Tirole (1988) has been replaced by price leadership: when prices reach the bottom, Red immediately increases his price, and Green soon follows suit.

The asymmetry in pricing and Buybox dominance criteria suggests that there might be meaningful differences between the sellers. There is presumably a reason why Amazon assigns the Buybox to Red when prices are tied. Indeed, there is a significant difference between Red and Green: Green is a third-party seller, and Red is Amazon itself. Nevertheless, the Buybox is roughly equally shared: Green has a Buybox share of 52.19%. Perhaps the platform realizes that sharing the Buybox is the best way to ensure that many sellers use the platform.

Given that the Buybox is essentially allocated to the cheaper merchant, it is easy to see why neither of the sellers achieves a high Buybox share: they have similar reaction times. In particular, the average reaction time for Green is 197s, and the average reaction time for Red is 247s. The average cycle takes 1.67 hours and has an amplitude of $0.10. While this might seem small, the cycling itself indicates the presence of potentially collusive profits: the welfare loss we are concerned with is not that of price volatility but that of high prices in general.

In this vein, we can focus on Green (a seller employing the repricing company's services) and exploit the cost information in our dataset as well as publicly available information on Amazon fees to work out his potential profit margin at the average price in the cycle[5]: 35%. This margin might not seem extraordinary, but note that the seller is directly competing with Amazon for this exact product. Thus, there is (i) perfect competition (as both sellers sell the same shoes in the same brand, size, and color), and (ii) we can expect costs to be bounded above by Green's costs for both sellers (as we would expect Amazon to be able to buy in bulk). Economic theory would thus suggest that profit should be very small (or zero when considering the seller's labor value).

[5] The net profit per sale is $34.21.

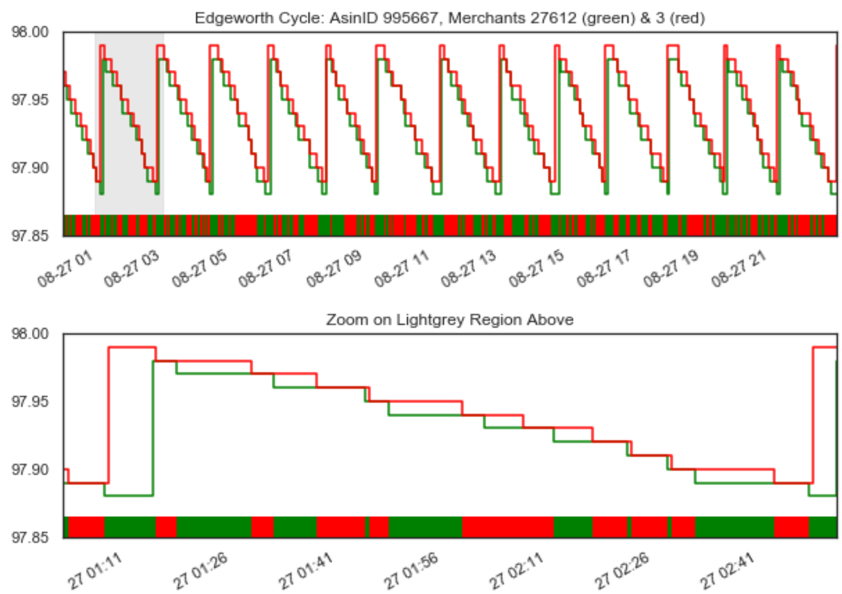


Figure 2: A Typical Cycle in the Pricing Data.

*Notes*: The upper panel depicts the prices of two different offers (green and red) for the same product. Prices decrease slowly over time before jumping back up. The lower panel zooms in on the light grey region in the upper panel. The bottom of each graph shows which offer has ownership of the Buybox at a given time.

# 3 Descriptive Evidence

## 3.1 Description of Data

This paper employs proprietary data from a repricing company that offers its services to Amazon third-party sellers. The company manages offer listings on its customers' behalf, allowing it to register with Amazon to receive price-change notifications. Notifications are sent within seconds if any of the offers on a given product change in price. They contain the price and some other information about the twenty lowest-priced offers on the product.[6]

Our main data source is the near-complete[7] set of 1,331,657,536 notifications received by the repricing company between 08/26/2018 and 03/25/2020. These notifications cover[8] 859,823 products, 159,112 merchants and 8,254,788 merchant-product pairs (i.e., offers). We further observe sales information (for the repricing company's clients only), self-reported costs, and a panel detailing which offers were repriced according to which rule when.

As Noel (2008) highlights, "in most cases where cycles are newly found, it is because finer and

[6]Offers are ranked by 'landed price' (i.e., price + shipping cost), and ties are broken randomly.

[7]The company data indicates that it received no notifications on 21/09/18 and fewer notifications than usual on 09/05/18.

[8]Though the data available to us is extensive, we emphasize that our empirical strategies below often drop a lot of irrelevant data in the service of our identifying assumptions – e.g., the RD can only utilize 14 842 offers for which we observe a transition from non-repricing to repricing.

newly available data reveals previously hidden cycles." Our data is attractive from this perspective as it samples prices at an infinite time resolution. Furthermore, as offers are grouped by product, we expect them to have high own-price elasticities. Also, many Amazon merchants are resellers for whom the assumption of a constant unit cost is an acceptable approximation. Finally, the Amazon data allows us to observe the price monitoring technology. As confirmed by conversations with the repricing company, it is precisely the price-change notifications we observe that the company (and supposedly others like it) uses to calculate and update the offers of their customers. Thus, we observe the complete information set that repricing strategies condition on.

There are also drawbacks to the data we employ. In particular, both the original Maskin and Tirole (1988) model and the equilibrium in delegated strategies we discuss below assume a duopoly – but the median product we observe has 8 offers. However, the number of *competitive* offers is typically far lower; and to the extent that additional offers are employing a "passive" constant-price strategy, they simply present a ceiling above which no algorithm can price.

## 3.2 Repricing Interfaces

Merchants can choose from a plethora of potential repricing companies, which try to differentiate themselves by, e.g., offering strategies "developed with game theory in mind" (SellerSnap) or "powered by machine learning" (Aura). However, a closer examination of repricing interfaces suggests that this wide variety in marketing claims is only partially reflected in the implemented strategies' sophistication. Finally, there is a small but burgeoning industry of sales-driven repricing. As this feature is rare (e.g., our data source has no such feature), we will ignore it in this paper and focus on competition-based repricing rules instead.

We exhibit two standard repricing interfaces in Figure 5: between the two, the possible strategies are almost identical. While these may be an outlier in terms of similarity, they embody the general gist of most interfaces well: merchants choose (i) by how much to undercut their competitors and (ii) what to do when such undercutting would lead to their price dropping below a pre-specified minimum. In such a case, many (but not all) repricers allow increasing the price to some pre-specified maximum. Minimum and maximum prices are specified separately for each product, and more advanced repricers often attempt to attract merchants by offering them some way to calculate these values automatically.

Thus, although the presence of cycles in our dataset might lead one to conclude that repricing algorithms are playing the MPE strategies proposed in Maskin and Tirole (1988), the interfaces we observe are not consistent with this interpretation. In particular, the Maskin-Tirole strategies require merchants to switch from just undercutting their rival's price to pricing at marginal cost once prices have dropped far enough. Furthermore, while prices are at marginal cost, merchants are required to randomize between resetting prices or keeping them unchanged, hoping that their rival might be the one to reset prices. Both randomization and jumps, however, are not implementable using the repricing software interfaces.

While there is variation in interfaces and implementable strategies, most repricers fall into two

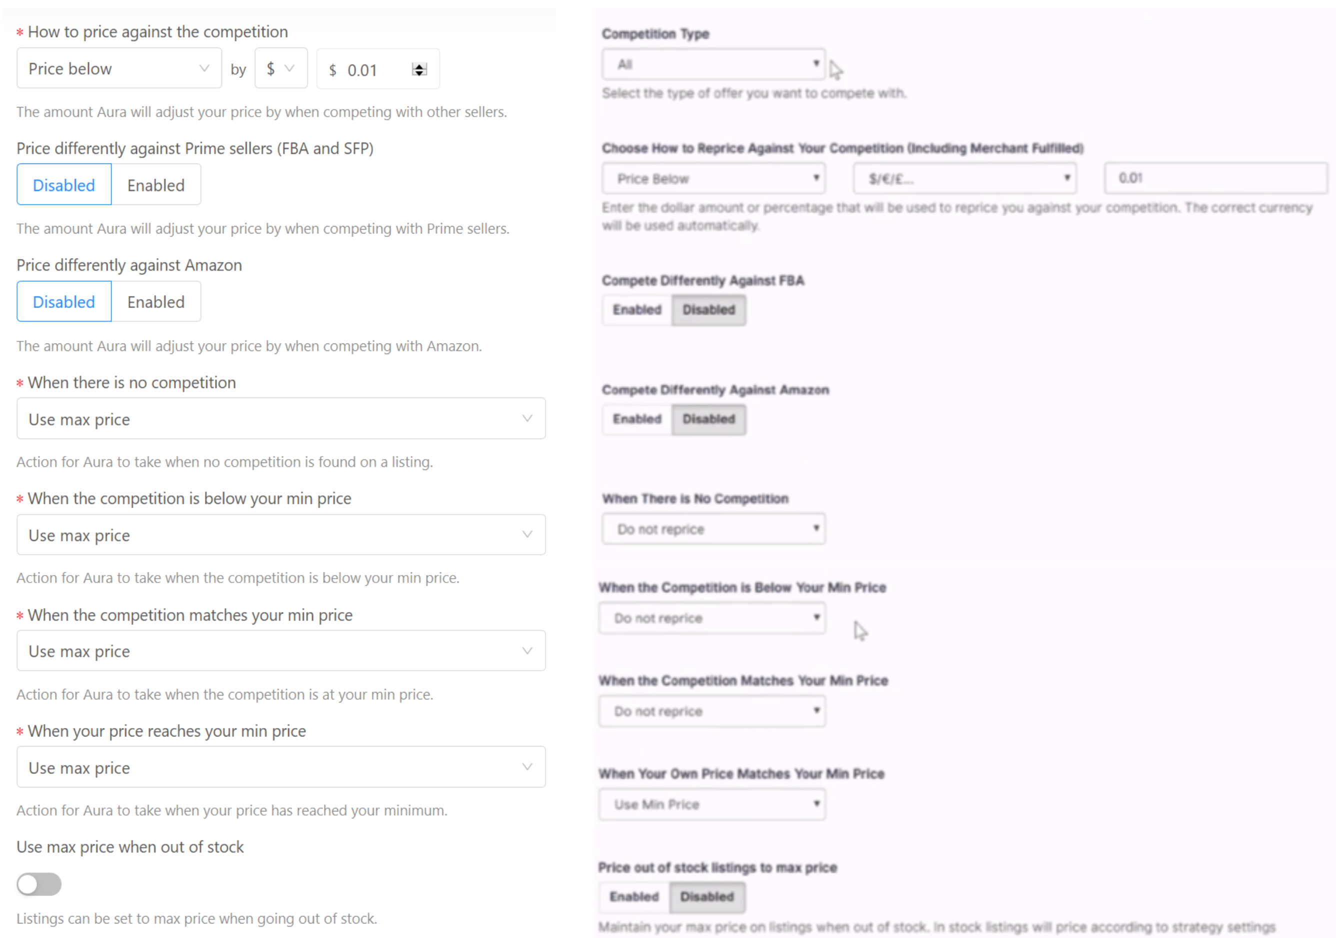


(a) Aura Interface (b) Informed Interface

Figure 3: Examples of Repricing Interfaces.

*Notes*: These pictures are screenshots of actual interfaces used by two repricing companies.

categories: they offer an undercutting or cycling strategy. Both strategies undercut the lowest-priced (relevant) competitor up to some minimum price. Once this minimum is reached, the cycling strategy will reset the price to some maximum, while the undercutting strategy will leave it unchanged.

### 3.3 Repricer Activation RD

When a merchant activates the repricer, he gives control of his prices to an algorithm. As a first step towards understanding how this may impact the marketplace, we are interested in the causal effect of activating the repricer on outcomes of the offer for which repricing was activated. To identify a causal effect, we exploit that signing up for the repricer and turning it on are separate actions. Some merchants activate the repricer for all their offers directly after signing up. Still, there is a delay for others, or they only start repricing for a subset of offers initially. This delay happens because (i) signing up is free, but repricing costs money (charged on a 'number of actively repricing offers' basis), and (ii) the repricing interface requires familiarization. We can thus perform a regression discontinuity design using time as the running variable and the activation of repricing as the cutoff.

Formally speaking, we estimate regressions of the form

$$y_{it} = \alpha_i + 1\{t \geq 0\} \times \beta + 1\{t < 0\} \times t \times \gamma_0 + 1\{t \geq 0\} \times t \times \gamma_1 + \epsilon_{it}, \tag{1}$$

where $i$ refers to an offer and $t$ refers to the number of days since the offer started automatically repricing. Here, $\alpha_i$ is an offer fixed-effect, $\beta$ is the coefficient of interest and the $\gamma$ coefficients allow outcomes to vary with the number of days since the repricer was activated. For each offer, we limit the sample to observations 'close' (in time) to the cutoff, i.e., $|t| < B$ where the bandwidth $B$ varies across specifications from $B = 19$ to $B = 5$. Out of an abundance of caution, we cluster the standard errors for all our results at the merchant level. As the date of repricer activation is correlated within merchant but far from perfectly so[9], our standard errors are hence conservative (Abadie et al., 2017).

The identifying assumption underlying (1) is that there are no *simultaneous* discontinuous changes in the various outcome variables on the date that the repricer is activated. This assumption would be violated if, e.g., merchants knew about an upcoming change in the Buybox algorithm and were expecting to lose Buybox share on a specific date, causing them to activate the repricing algorithm on precisely this date (and not one day earlier). In general, discontinuous changes in non-treatment covariates (like the parameters of the Buybox algorithm) seem plausibly related to merchant repricing behavior. However, it seems implausible that merchants anticipate *the exact date* of such changes or react to them within a day or two. Furthermore, to violate our identification strategy, the merchants would have to be already signed up with the repricing company, awaiting a discontinuous change in exogenous covariates only to react immediately by activating the repricer. This scenario seems implausible. Finally, our identifying assumption should be contrasted with the parallel trends assumption that would underlie a differences-in-differences approach. While such an approach yields similar estimates, the assumption of parallel trends is violated, perhaps because

[9] A regression of repricer activation date on merchant fixed effects yields $R^2 = 0.37$.

merchants familiarizing themselves with the repricing interface learn more about the pricing of their offers and start making small changes to their pricing manually.

We report the results of running our RD models in Table 1 and illustrate the widest-bandwidth implementation in Figure 4. In this figure, the lightly shaded area indicates the period for which we cannot ascertain treatment status due to sampling frequency restrictions. In the top-left panel, we see that offer-level prices decrease dramatically on impact by an average of 5.2%. As indicated by the bottom-left panel, this price decrease is accompanied by a significant increase in Buybox share of 5.3pp. The additional demand exposure offsets the lower prices; hence, we see in the bottom-right panel that profits increase by 39.7%. Profits increase because prices do not fall more than 'necessary,' i.e., the repricer seems to undercut the current Buybox price: in particular, on impact, Buybox prices do not move.

There are some strong post-trends in Figure 4. Therefore, we extend the post-window in Figure E.2 (see appendix) and find that (i) the own-offer price stabilizes in the long run, (ii) the Buybox price keeps falling for two months (but never reaches the extent of the decline in own price), (iii) the Buybox share stabilizes after a month and (iv) profits decline to the pre-repricing level after 1.5 months.

To ensure our results are robust, we follow the suggestions of Hausman and Rapson (2018), who note that a McCrary (2008) density test is not applicable when time is the running variable and suggests alternative robustness tests[10]. Table 1 shows that our estimates are robust to the bandwidth used. In the appendix, we also explore robustness to dropping observations to create a balanced panel (Table E.1), robustness to autoregression in the dependent variable (Table E.2) and robustness to a quadratic specification (Table E.3). The results are fundamentally unchanged, though the quadratic specification (as expected) requires a larger bandwidth to generate significant results, and the $AR(1)$ specification estimates smaller (but still statistically significant and directionally equivalent) treatment effects.

To summarize, Figure 4 tells us that the typical merchants' offers see their prices drop when repricing is activated: in particular, the repricer will typically try to target the highest price that allows the merchant to obtain the Buybox. As discussed above, we will refer to this kind of strategy as 'undercutting' (see below for a more formal definition). But, for now, note that while this sort of strategy is bound to perform very well against a marketplace populated by merchants not repricing themselves, it performs poorly against other agents employing the same strategy. Indeed, prices will race down to cost quickly when an undercutter faces another.

### 3.4 The Importance of Resetting Prices

To avoid price wars, merchants may want to increase prices periodically. Indeed, most repricing interfaces offer this ability by either letting merchants increase prices at regular periodicity or in reaction to a select number of events (e.g., when undercutting an opponent's price would lead to a price that lies below a pre-specified minimum). We refer to this behavior as resetting.

[10] Note that these alternative tests cannot evaluate the possibility of sorting into treatment.

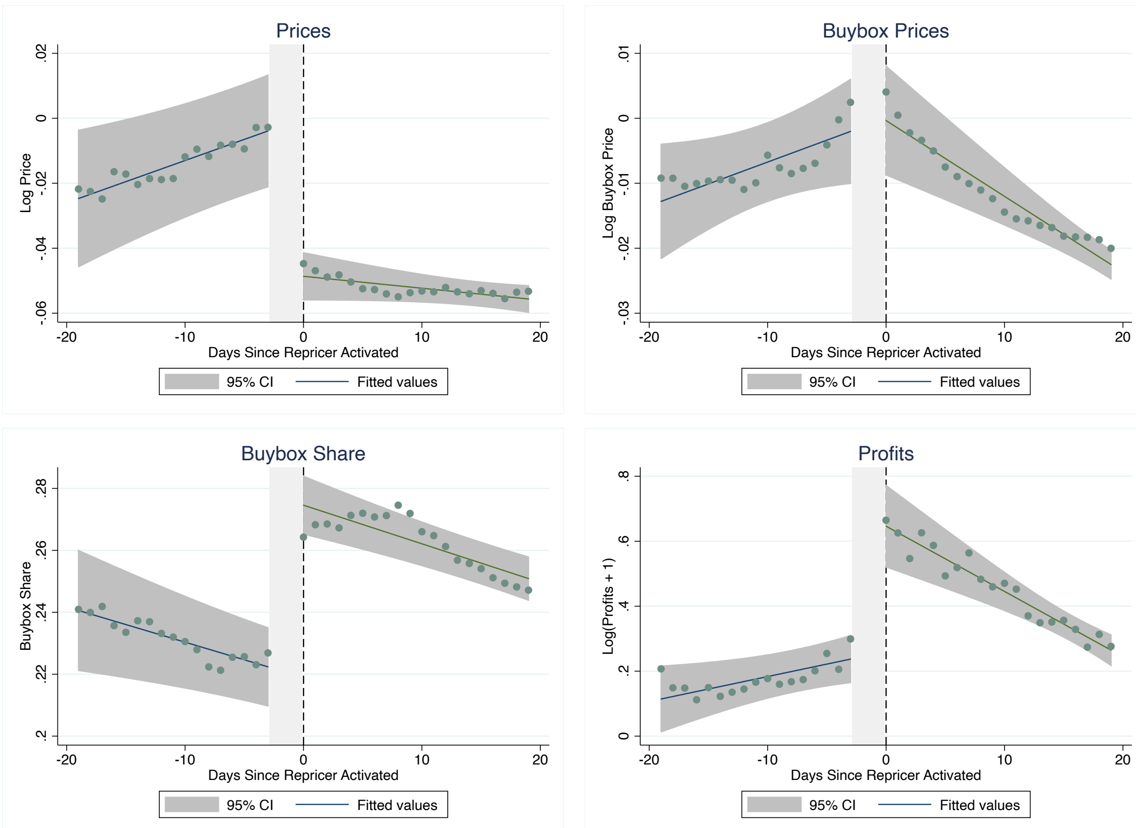


Figure 4: Activating the Repricer Lowers Prices and Increases Sales.

*Notes*: Results from RD around repricer activation using transition of 14 842 offers by 121 merchants from non-repricing to repricing for 'Prices' and 'Buybox Share.' Due to loss from merging in additional data, the 'Buybox Prices' results are based on 14 052 offers from 120 merchants and the 'Profits' results are based on 1781 offers from 81 merchants. The light-gray area indicates the time during which we cannot ascertain treatment status due to sampling frequency. The dark-gray area indicates the 95% confidence interval derived from standard errors that have been clustered at the merchant level. These confidence intervals are generally too conservative given the imperfect within-merchant correlation of treatment assignment (Abadie et al., 2017). See Figure E.1 for the distribution of included events over time.

| Dep. Var. \ Bandwidth | ±19 | ±15 | ±10 | ±5 |
|---|---|---|---|---|
| Log Offer Price | -0.052*** | -0.053*** | -0.050*** | -0.053*** |
| | (0.010) | (0.010) | (0.011) | (0.008) |
| Log Buybox Price | -0.003 | -0.004 | -0.005 | -0.009** |
| | (0.003) | (0.003) | (0.004) | (0.003) |
| Buybox Share | 0.053*** | 0.050*** | 0.043*** | 0.037*** |
| | (0.011) | (0.010) | (0.012) | (0.013) |
| Log(Profits + 1) | 0.397*** | 0.368*** | 0.322*** | 0.249** |
| | (0.076) | (0.075) | (0.073) | (0.116) |

Table 1: RD Estimates Are Robust to Bandwidth.

*Notes*: This table provides $\hat{\beta}$, the estimated effect of repricer activation on various outcome variables (row) and for different estimation windows (column). As our specification is linear, these results can be interpreted as the results from a local linear regression with a rectangular kernel of bandwidth given by the column title. For instance, column "±15," contains the results of running regression (1) on a sample that has been limited to include observations from each offer only 15 days before and after the repricer was activated for said offer. Standard errors clustered at the merchant level in parentheses. *** p<0.01, ** p<0.05, * p<0.1

#### 3.4.1 Does Resetting Work?

For resetting to be effective, the opponent must react to the resetting merchant's price increase with a price increase of their own. While this may seem to follow trivially from an undercutting strategy, there are nevertheless reasons we may worry about the success of such a strategy: e.g., it relies on the opponent's ability to monitor prices and respond quickly enough.

To determine whether resetting raises prices, we need to disentangle selection and treatment effects. As resetting is designed to counter opponents that employ an undercutting strategy, merchants are presumably more likely to activate resetting strategies on products where they face tough competition. To avoid conflating these effects, we exploit the fact that the repricing system at our source company sometimes gets overwhelmed and cannot execute all intended repricings. In such a case, the company employs a queue system. While the position in the queue is not random, it is determined by the exact time a merchant intended prices to be reset, which is plausibly exogenous. Hence, we consider the repricing queue a natural experiment and utilize the resulting quasi-random variation to identify the treatment effect of resetting.

We restrict attention to the set of offers for which resetting was activated and run the following regression

$$y_{it} = \alpha_i + 1\{Reset\}_{it} \times \beta + \epsilon_{it},$$

where $i$ refers to an offer and $t$ to a date. Here, $y_{it}$ is one of several outcome variables, $\alpha_i$ is an offer fixed-effect and $1\{Reset\}_{it}$ is an indicator for whether or not a scheduled reset actually took place. Our identifying assumption that $\mathbb{E}[1\{Reset\}_{it}\epsilon_{it}] = 0$ is plausible in the subset of offers and dates for which resetting was configured to happen. For it to be violated, there would have to be a

| | Dependent Variable | Reset Happened? | | N | $R^2$ | Offer FE |
|---|---|---|---|---|---|---|
| (1) | Own # Price Changes | 1.592*** | (0.297) | 24,985 | 0.430 | Yes |
| (2) | Competitor # Price Changes | 0.269*** | (0.0814) | 157,163 | 0.440 | Yes |
| (3) | Own Log Price | 0.0774*** | (0.00995) | 24,985 | 0.979 | Yes |
| (4) | Competitor Log Price | 0.0123*** | (0.00411) | 157,163 | 0.988 | Yes |
| (5) | Buybox Log Price | 0.0194*** | (0.00678) | 23,784 | 0.982 | Yes |

Table 2: Price Resets Are Effective At Raising Opponent Prices.

*Notes*: This table shows the results of the regression of various dependent variables (one per row) on an indicator for whether a price reset happened in the preceding day, using a sample of 831 products for which a price reset was configured but not always executed. Standard errors are clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1

connection between the repricing queue becoming overwhelmed for a particular offer and the prices of merchants competing with it. In particular, as there is a single repricing queue and offers for many unrelated products are repriced in that same queue, we believe it unlikely that a reset is not executed because of a flurry of repricing activity on one particular offer (which could be endogenous). Instead, unexpected activity on *other* products will be the primary cause for configured resets that are not executed.

Our results are in Table 2, where we separate the effects of resetting on own, competitor, and market prices. To begin with, however, we verify in the first row that the resetting merchant has more price changes on days we classify him as resetting. Similarly, the effect of a reset on own price is positive as well as statistically and economically significant in the third row: on days when a reset occurs, the resetting offer's price is on average 7.74% higher than on days when no reset occurs. In an IV framework, this suggests a solid first stage. The main results of interest are in the fourth and fifth rows. Here, we analyze the impact of one merchant's resetting on the average competitors' price and the Buybox price, respectively. The average competitor price increases by 1.23%, and the Buybox price – our proxy for the market price – increases by 1.94%. There is thus a strong indication that resetting has the desired effect of coaxing competitors to raise their prices.

We perform some robustness checks for our results in Appendix F. First, we verify that the results are driven by serious competitors (defined as competitors that, on average, don't price much above the resetting merchant) in Table F.1; indeed, the responses of serious competitors are almost twice as strong. Furthermore, we control for configured resets with a dummy variable (rather than conditioning on products for which a reset was configured) in Table F.2; the table suggests that the selection effect is negative (as expected). Finally, we treat the execution of resets (in the subsample for which resets are configured) as an instrument for estimating the effect of the repricing merchant's price on competitor prices in Table F.3; consistent with our previous results, we find a pass-through of approximately one-sixth.

### 3.4.2 Is Resetting Strategic?

Having established that resetting can successfully raise prices, we would like to know whether merchants think strategically about when to increase prices. In Appendix D, we develop an algorithm that can identify cycling offers and when they are reset. We now present some descriptive and narrative evidence that there is at least some general awareness that the exact nature of the cycle played will influence payoffs. This evidence relies on a critical difference between the theoretical model in Maskin and Tirole (1988) and the data we observe: heterogeneity of sales across time. In particular, sales generally happen when people are awake. The fact that night-time sales are not common has strategic implications for the sellers: when sales probabilities are low, the costs of 'resetting' (i.e., increasing) prices are also low. Their promotional material confirms that repricing services are aware of this: e.g., one company promises that it "will allow you to stop automated repricing for a period every day (normally the late evening and early morning when sales are lowest) and reset to maximum (if preferred), as this can often help drive prices back up across all competition" (repricerexpress.com).

Given this awareness, it is not surprising to see that resets of daily cycles do indeed occur during the night hours when sales probabilities are lowest. Figure 5a depicts histograms of sales (in blue, measured from the top) and reset times[11] (in orange, measured from the bottom) of approximately day-long cycles by the hour of the day. Sellers seem to reset day-long cycles almost exclusively between 1 am and 9 am (Chicago time), with virtually all of the mass in an hour-long window around 4 am.

As becomes apparent when looking at the typical day-long cycle depicted in Figure 5b, the distinct pattern of reset timing is not the only thing special about day-long cycles. These cycles also feature an 'excessive' amplitude: the first price decrease after the reset is much larger than subsequent price decreases. We expect the decrease to be the same at each (decreasing) step of an Edgeworth cycle. Instead, we see a significant decrease at the first step followed by much smaller decreases.

The excessive nightly resets may be rational even though they do not yield any benefit on the equilibrium path. This is because nightly resets are essentially free from the perspective of lost sales. At the same time, if equilibrium is perturbed, there might be a substantial benefit to briefly raising the price by more than necessary. Suppose a rival had also configured a maximum price, for instance. In that case, a merchant might want a strategy that robustly raises his price to his maximum without wanting to adjust his strategy every time his rival adjusts his maximum price.

## 3.5 Summary Statistics

We discuss how the likelihood of cycling varies with the product category in Appendix A. For products with at least one cycle recognized, Figure 6a depicts the distribution of the average cycle length in hours. As discussed above, many products' cycles take 24h on average.

[11] I define the "reset time" as the midpoint between the time at which prices are raised and the time at which the following reaction occurs.

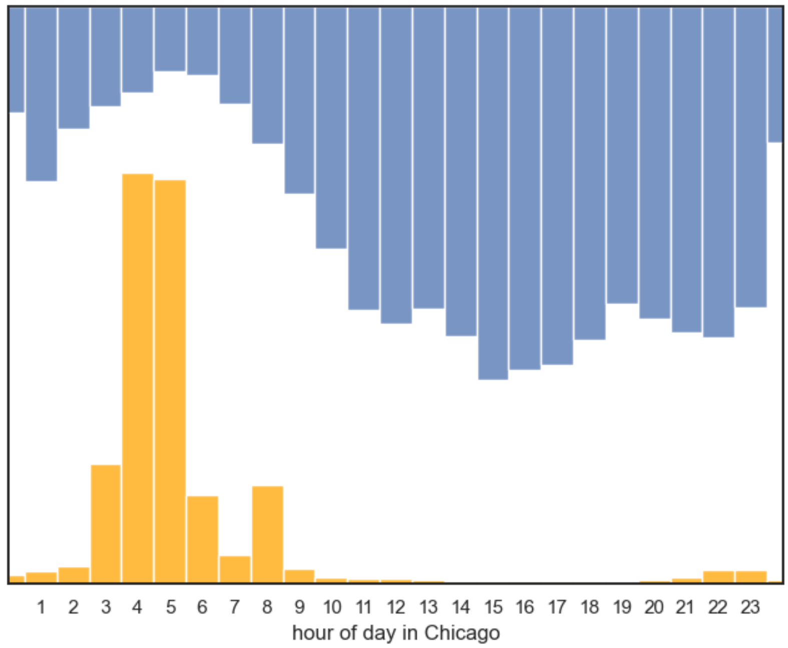


(a) The Reset Times of Daily Cycles

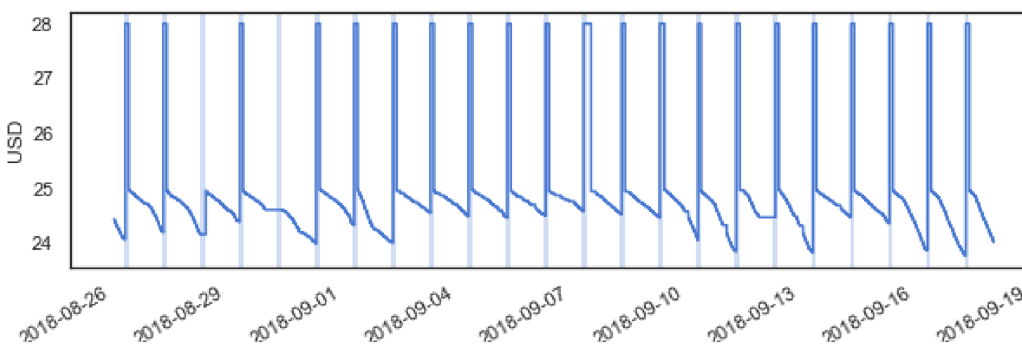


(b) Typical Day-Long Cycle

Figure 5: Daily Cycles.

*Notes*: The top figure shows the distribution of the reset times of 434,438 day-long cycles (in orange; measured from the bottom) and the distribution of 9,687,301 sales (in blue; measured from the top) by the hour of the day. We define a day-long cycle as a cycle that takes 24h $\pm 1h$ from peak to peak. The top figure illustrates that day-long cycles are almost exclusively reset at night when the overall sales activity is lowest. The bottom plot shows the price (on the vertical axis) of a typical offer that is cycling daily against the date and time (on the horizontal axis). The shaded regions correspond to 2 am to 5am Chicago time, and we see that during those times, the offer's price increases dramatically, only to be lowered again significantly when the next day begins.

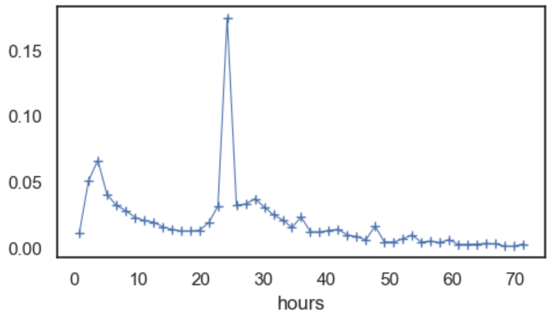


(a) Distribution of Cycling Period

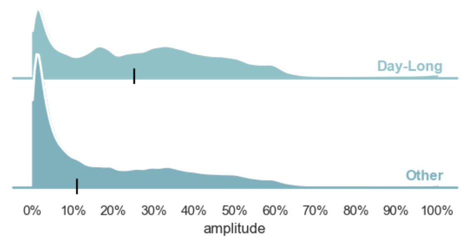


(b) Distribution of Cycling Amplitude

Figure 6: Distributions of Cycle Summary Statistics.

*Notes*: The left panel shows the distribution of average cycling periods (i.e., the length of time between resets) for all products. The right panel shows the distribution (black bars: median) of cycle amplitude as a fraction of the maximum price separately for day-long cycles and other periods.

We plot the distribution of the average amplitude as a fraction of the mean price of the offer in Figure 6b. We separate day-long cycles, which have substantially larger amplitudes: the median (cycling) product has an average amplitude that is 25.21% of the mean price. While other cycles have smaller amplitudes, the median for them is still 11.06% of mean price. These amplitudes are similar to those documented for other industries in the literature. For comparison, Noel (2018) states that the typical amplitude of cycles found in gasoline markets is "about ten percent of the price" which is also consistent with the 13% amplitude found in Noel (2007).

## 4 Equilibrium in Delegated Strategies

We will model sellers as choosing between two active repricing strategies ($U$ for undercutting and $C$ for cycling) and a passive strategy $F$. Our discussion will assume that the market starts at $t = 0$ with all agents choosing prices according to $F$ – this is meant to resemble the situation before Amazon introduced the API that made repricing possible.

Every period, one agent is allowed to update her chosen repricing technology. Agents choose myopically, i.e., they do not anticipate the future evolution of the marketplace. Furthermore, to mirror the heterogeneity in repricing strategies observed in the data, agents face costs related to the complexity of the repricing technology they choose. In particular, agents can select one of three technologies, each of which maps a sole competitor's price $p_j$ into the merchant's price, potentially as a function of parameters $(\underline{p}, \overline{p})$ to be chosen by the merchant:

$$r_F(p_j) \sim F(\cdot), \quad r_U(p_j; \underline{p}) = \max\{p_j - k, \underline{p}\}, \quad r_C(p_j; \overline{p}, \underline{p}) = \begin{cases} p_j - k & \text{if } p_j - k \in (\underline{p}, \overline{p}], \\ \overline{p} & \text{otherwise.} \end{cases}$$

Here $k$ is the minimum unit of currency and the minimum amount by which the price can be decreased. We assume that all agents can access $r_F$ for free while the other technologies are available at costs $c_U$ and $c_C$. More realistically, both repricing technologies' costs should be declining over time. In particular, Amazon rolled out free access to their own 'Automate Pricing' tool (effectively a repricing algorithm of type $U$) in 2016 and thus effectively set $c_U = 0$.

Each period after making (or not making) their technology choice, agents are randomly matched in pairs to list offers on the same products. At this stage, they are committed to their technology but not to its parameters: in particular, they play a Nash Equilibrium in parameter choices. Thus, before discussing the dynamic evolution of repricing technology shares, we must examine the within-period equilibria. For brevity's sake, we focus on the pairing $(U, C)$ in the main text; all other pairings are discussed in Appendix C as they do not lead to supra-competitive profits.[12]

## 4.1 Within-Period Equilibria

To discuss potential equilibria from the pairing $(U, C)$ and highlight the difference with the Maskin-Tirole strategies, we begin by considering a slight modification of their original example. There are two merchants who both produce output at marginal cost $c = 0$ and face industry demand $D(p) = 20 - p$. As a first-order approximation of the 'Buybox' mechanism, we assume[13] that only the currently lower-priced seller is exposed to this industry demand; the other seller faces zero demand. Prices must lie on a grid $p \in \mathbb{N}$ (i.e., $k = 1$) and can only be changed in alternating periods: one merchant may change prices in odd, and the other may only change prices in even periods.

In Maskin-Tirole, agents' strategies are response functions $r_i(p_j)$, which specify the own price as a function of the last price set by the opponent. In our setup, however, agents' response functions are parameterized, and their strategies are the parameters of the functions; once these parameters are chosen, the response function is implemented mechanically. In line with our discussion above, the undercutting merchant determines a minimum price $\underline{p}_U$, and the (potentially) cycling merchant selects both a minimum $\underline{p}_C$ and a maximum price $\overline{p}_C$.

We can perform a simple grid search to confirm that the[14] equilibrium parameter choices of the merchants are

$$\Big((\underline{p}_C, \overline{p}_C), \underline{p}_U\Big) = \Big((16, 4), 0\Big).$$

Furthermore, following Maskin and Tirole (1988) we can find an Edgeworth MPE in this setup[15]. Figure 7 displays simulated sample paths from our delegated strategy equilibrium (left panel) and an Edgeworth MPE (right panel). While the delegated strategies produce a similar cycling phenomenon,

[12]The reader may have noticed immediately that $(U, U)$ pairings lead to pricing at cost; it can be shown that this also is the case for $(C, C)$ pairings.

[13]To be clear, this assumption is wrong in the sense that we know (i) not all consumers purchase through the Buybox and (ii) the Buybox is not infinitely price elastic but has a price elasticity of around -19 (Lee and Musolff, 2021). Nevertheless, the assumption serves as a good approximation, and from interviews with employees of the repricing company, sellers seem to intuitively utilize a heuristic similar to that embedded in this assumption for assessing how much demand their offer will be exposed to.

[14]Technically speaking, $U$ has multiple payoff-equivalent equilibrium choices for $\underline{p}_U$.

[15]There are typically multiple Edgeworth MPEs, but they all exhibit the features we discuss. For illustration purposes, we chose an MPE in which resets do not take too long to occur.

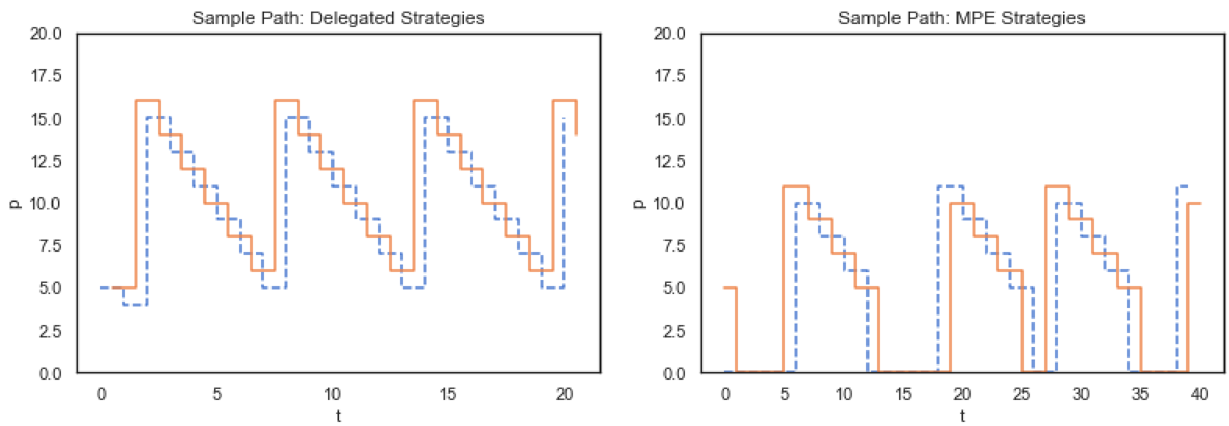


Figure 7: Sample Paths under MPE vs Delegated Strategies.
*Notes*: This figure provides example price paths for cycling under delegated strategies (left panel) and Markov-perfect strategies (i.e., the Maskin-Tirole equilibrium).

there are two key differences. Firstly, the price never reaches marginal cost in the delegated strategy sample path. Secondly, there is no war of attrition at the bottom of the cycle: instead, it is always the same merchant who leads the cycle back up by resetting prices.

These differences are crucially related to the restricted strategy space that our delegated strategy model imposes. Given the restrictions in the repricing interface, there is simply no way for merchants to jump down to costs. Thus, the agents cannot punish opponents for undercutting the cycle; indeed, the unique equilibrium amongst C agents is to price at cost. However, as long as the opponent plays a simple undercutting strategy, no punishment is necessary, and a cycle can be sustained.

Given that the price path under delegated strategies avoids lengthy periods of marginal cost pricing, it is natural to suspect that it outperforms the MPE in terms of profits. Indeed, with discount factor $\delta = 0.99$ we find that while the merchants' joint profit under MPEs is 50% of monopoly profit[16], under delegated strategies this figure is 88%.

While this is already alarming, we will now see that the closeness of profits to monopoly profits understates the welfare loss. To this end, we build an analytic model of the parameter choice game. Moreover, we will make a simplifying assumption that renders the model tractable: we assume that prices are moving smoothly during the decreasing phase of the cycle. This is not true in the real world: the minimum unit of currency $k$ ensures that jumps characterize real price paths. But as these jumps are small, the assumption of smooth price paths will be a good approximation for sufficiently small $k$.

However, we cannot simply let $k \to 0$ to arrive at our continuous approximation: doing so would distort the incentives of the resetting agent by making resets costless in the limit. To prevent this, we reinterpret $k$ as the (fixed) rate at which prices decrease during an undercutting phase and instead let the time interval $\Delta$ that passes between two subsequent undercutting steps tend to zero. Thus,

[16]Note that Maskin and Tirole show that as $\delta \to 1$, in symmetric Edgeworth MPEs both firms must earn at least 50% of monopoly profits.

repricing agents undercut each other more and more often but by smaller and smaller amounts. Furthermore, we introduce a time penalty for resetting prices. After increasing prices above his rival, a merchant must wait one complete unit of time until he is allowed to reprice again.

The above discussion implies that the indirect utility from a cycle with trough $\ell$ and peak $u$ to a resetting (non-resetting) agent is given by $V_r(\ell, u)$ ($V_{nr}(\ell, u)$), where

$$\begin{aligned} V_r(\ell, u) &= \frac{1}{2(u-\ell+k)} \int_\ell^u \pi(p)dp, \\ V_{nr}(\ell, u) &= \frac{1}{2(u-\ell+k)} \int_\ell^u \pi(p)dp + \frac{k}{u-\ell+k}\pi(u). \end{aligned}$$

Here, $\pi(\cdot)$ maps a price $p$ into the profit a monopolist setting this price would obtain. Note how the parameter choices translate into the payoffs to the agents. Assuming $\overline{p}_C > \underline{p}_C > \underline{p}_U$ (true in equilibrium), the price path will be a cycle from $u = \overline{p}_C$ to $\ell = \underline{p}_C$ that is reset by C. Furthermore, the cycle is not influenced by $\underline{p}_U$ other than that $\underline{p}_U$ needs to be sufficiently low such that $C$ does not have an incentive to abandon the cycle and set a constant price just below $\underline{p}_U$.

**Proposition 1.** *Let $(\ell^*, u^*) := \arg\max V_r(\ell, u)$ and $x^* := \min\{x : \pi(x) = V_r(\ell^*, u^*)\}$. When a $U$ agent faces a $C$ agent, there is a continuum of equilibria given by*

$$\left\{ \left( (\underline{p}_U), (\underline{p}_C, \overline{p}_C) \right) : \underline{p}_U \le x^*, \underline{p}_C = \ell^*, \overline{p}_C = u^* \right\}.$$

*These equilibria all result in identical price paths and payoffs.*

(All proofs in Appendix B.) Hence, in equilibrium, $C$ can choose his preferred cycle. Which cycle he prefers is determined by the trade-off between wanting to reset as seldom as possible and not letting prices drift too far from the monopoly price. In particular, the shorter the cycle, the closer profits in most periods stay to monopoly. However, the longer the cycle, the less frequently it has to be reset (which is costly as it requires a period of zero profits).

**Proposition 2.** *If $\pi(\cdot)$ is single-peaked, the solutions to $\arg\max V_r(\ell, u)$ must satisfy the FOCs.*

Thus, the optimal cycle solves:

$$\pi(u) = \frac{1}{u-\ell+k} \int_\ell^u \pi(p)dp, \tag{2}$$

$$\pi(\ell) = \frac{1}{u-\ell+k} \int_\ell^u \pi(p)dp. \tag{3}$$

Both (2) and (3) equate the marginal benefit of extending the cycle (LHS) to the marginal cost of extending the cycle (RHS). The tradeoff for (3) is illustrated in Figure 8 where the marginal benefit is shaded green and the marginal cost red (after eliminating the 'double counted' section). From the FOCs, we immediately have the following two results.

**Proposition 3.** *If $\pi$ is symmetric around the monopoly price $p_m$ and $\tilde{\pi} = g \circ \pi$ where $g'' \ge 0, g' > 0$ and $g(0) = 0$, then $\tilde{\ell}^* \ge \ell^*$ and $\tilde{u}^* \le u^*$.*

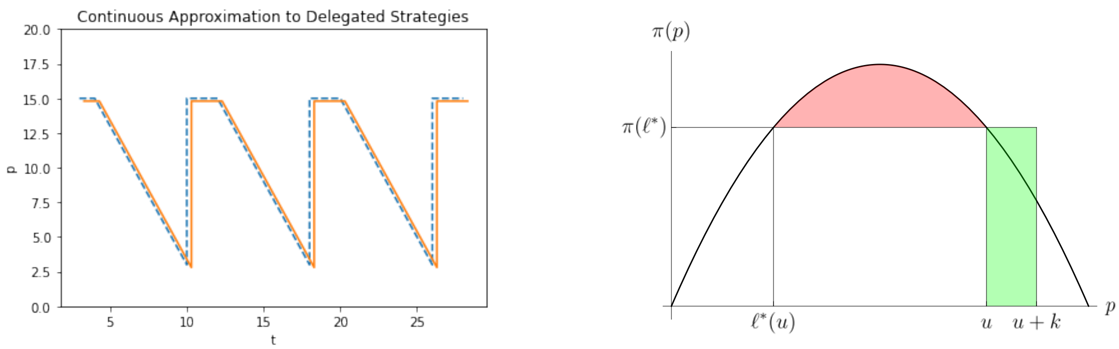


Figure 8: The Continuous Approximation (left) & The FOC for Choice of $\ell$ (right).
*Notes*: This figure illustrates our continuous approximation to the discrete pricing path generated by delegated strategies. The left figure shows the price path that our continuous approximation implies (see LHS of Figure 7 for the path being approximated). The right figure shows the cycling merchant's tradeoff when deciding whether to extend the cycle by lowering the price $\ell$ at which a reset is initiated. The marginal benefit is shaded green, and the marginal cost red. For instance, an increase in $k$ would increase the green region, leading to a lower choice of $\ell$, i.e., a longer cycle.

**Proposition 4.** *If $\pi(\cdot)$ is symmetric around the monopoly price $p_m$, the average price during the decreasing phase of the cycle will satisfy $p_{avg} = p_m$.*

Thus, the situation is worse than indicated in the Maskin-Tirole comparison example above. In equilibrium between $C$ and $U$, prices will spend at least half of the time *above* the monopoly price. From a static perspective, this time is a Pareto loss to the economy: all merchants and consumers would be better off if prices were magically lowered. However, regularly exceeding the monopoly price is necessary for the tacitly collusive scheme between $C$ and $U$ in a dynamic setting.

## 4.2 The Evolution of the Repricing Economy

We can think of the game of repricer choice as a repeated stage game. At each stage, firms are randomly paired and play the stage game implied by their repricing technology. For instance, if a $U$ and $C$ firm meet, they play the equilibrium we just discussed; the equilibria for other pairings are in Appendix C. Plugging in the equilibrium parameter choices, the stage game of technology choice then has payoffs given by Table 3, where

$$V_{ff} = \int \pi(p)[1 - F(p)]dF(p), \quad V_f = \int \pi(p)dF(p), \quad V_{nr}^* = V_{nr}(\ell^*, u^*), \quad \text{and} \quad V_r^* = V_r(\ell^*, u^*),$$

are the within-period payoffs. While $V_{nr}^*$ and $V_r^*$ follow from our discussion above, $V_f$ refers to the payoffs that $U$- and $C$-types obtain when facing an $F$-type. $V_{ff}$ is the payoff that an $F$-type obtains from facing another $F$-type.

We further introduce $c_U > 0$ and $c_C > 0$ to measure the costs associated with choosing more complex repricing strategies. Clearly $V_f > V_{ff}$ but we further put a restriction on costs such that

| | | Player 2 | | |
|---|---|---|---|---|
| | | $F$ | $U$ | $C$ |
| Player 1 | $F$ | $(V_{ff}, V_{ff})$ | $(0, V_f - c_U)$ | $(0, V_f - c_C)$ |
| | $U$ | $(V_f - c_U, 0)$ | $(-c_U, -c_U)$ | $(V_{nr}^* - c_U, V_r^* - c_C)$ |
| | $C$ | $(V_f - c_C, 0)$ | $(V_r^* - c_C, V_{nr}^* - c_U)$ | $(-c_C, -c_C)$ |

Table 3: Repricer Choice Stage Game Payoffs.

*Notes*: This table provides the payoffs to the repricer choice stage game: e.g., if P1 has a fixed price strategy (F) and P2 has an undercutting strategy (U), the payoff to P1 will be 0 and the payoff to P2 will be $V_f - c_U$.

$V_f - c_U > V_{ff}$. Furthermore, we will assume that for all $\alpha \in [0, 1]$,

$$\max\{\alpha V_r^* - c_C, (1-\alpha)V_{nr}^* - c_U\} > 0$$

so that it is always better to employ some type of repricer rather than just choose random prices. Note that this implies that $F$ is never a best response. The unique *symmetric* Nash equilibrium of the game is a mixed-strategy equilibrium: both players play U with probability

$$\gamma = \frac{V_{nr}^* + c_C - c_U}{V_{nr}^* + V_r^*},$$

and $C$ with probability $1 - \gamma$.

We now analyze the dynamic evolution of the repricing choices in an evolutionary model. In particular, we follow the model for the dynamic evolution of play in general games proposed by Kandori et al. (1993) and generalized in Kandori and Rob (1995). Their key assumptions are threefold: firstly, there is inertia, i.e., agents only adjust to best responses slowly. Secondly, agents are fundamentally myopic in that they best respond to the *current* distribution of strategies in the population but do not anticipate the future evolution of strategy shares. Finally, there is a small probability that agents make a (random) mistake.

Given these assumptions, the dynamic evolution of play can be understood as a Markov process. The associated state space consists of triplets $(n_F, n_U, n_C)$ describing how many players are playing each of the three strategies. For a given error rate $\epsilon$, denote the corresponding Markov transition matrix by $P(\epsilon)$. As KMR emphasize, $P(\epsilon)$ is aperiodic and irreducible. Hence, we know from standard Markov theory that there exists a unique stationary distribution $\mu(\epsilon)$ such that $\mu(\epsilon)P(\epsilon) = \mu(\epsilon)$. Furthermore, for any initial distribution $q$ over the state space, we have that $\lim_{t\to\infty} qP(\epsilon)^t = \mu(\epsilon)$, i.e. the stochastic process will converge to the stationary distribution from any starting point.

With a positive error rate $\epsilon$, the stationary distribution $\mu(\epsilon)$ will put a positive probability on all possible states. As the error rate tends to zero, KMR show that the stationary distribution converges to a unique limit $\mu^*$. Returning to the game described in Table 3, we have the following trivial implication of KMR's Theorem 5:

**Proposition 5.** *Suppose that in each period only one player is allowed to adjust their strategy.*

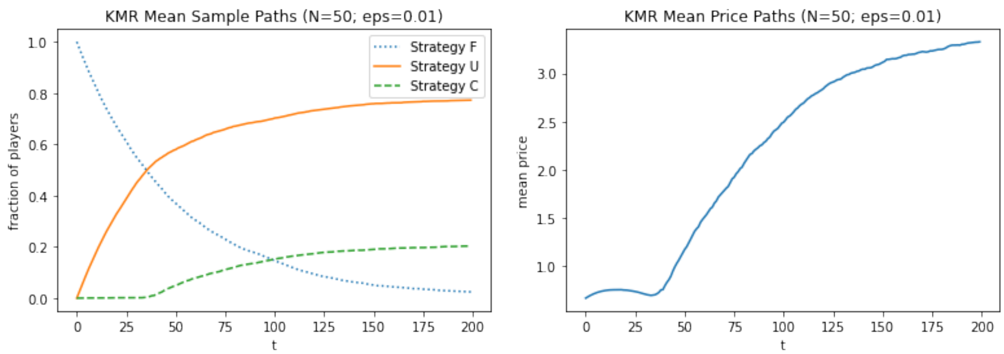


Figure 9: The Third-Party Marketplace Adjusts After Introduction of Repricers.
*Notes*: This figure illustrates the mean adoption paths of the three possible repricing strategies (left panel) and the resulting path of mean prices.

*Then the limit distribution puts probability 1/2 on* $(0, \lfloor \gamma N \rfloor, N - \lfloor \gamma N \rfloor)$ *and probability 1/2 on* $(0, \lceil \gamma N \rceil, N - \lceil \gamma N \rceil)$.

**Corollary 1.** *The mean price under the limit distribution is given by*

$$p_\infty = [1 - 2\gamma(1-\gamma)]c + 2\gamma(1-\gamma)p_m,$$

*where* $p_m$ *is the monopoly price and* $c$ *is the unit cost.*

Thus, no matter at which repricing strategy shares we start, the economy will eventually converge to a mix of $U$ and $C$ agents. This mix naturally pins down the mean price. Typically (e.g., if agents were close to pricing at Bertrand-Nash before the introduction of repricers), the mean price at the stationary distribution will be higher than the initial mean price. However, the extent to which prices increase is limited by a coordination problem: given the current strategy space, cycling can only be achieved by pairs of undercutters and cyclers.

Staying with the evolution of the mean price for the moment, there is another interesting and perhaps surprising result: the transitional dynamics are far from monotone.

**Proposition 6.** *If* $p_\infty > p_0$ *and* $c_C - c_U$ *is sufficiently high relative to the dispersion of the fixed pricers' price distribution* $F(\cdot)$*, then prices will first increase, then decrease and then increase again. Furthermore, prices will dip below the initial price.*

Thus, even if the mean price is currently *decreasing* due to the introduction of repricers into the marketplace, such a decrease might merely pave the way for future price increases.

### 4.3 Simulation

To illustrate the evolution of the repricing economy, we simulate the KMR Markov process. In particular, we fix $c_U = 5$, $c_C = 27$, $F(p) \sim U[0, 2]$ and (as above) let $\pi(p) = (20 - p)p$. Then it can be verified that $V_{ff} \approx 6.33$, $V_f \approx 18.66$, $V_r^* \approx 42.07$ and $V_{nr}^* \approx 51.46$. We fill these values in the payoff matrix above, set the error probability to $\epsilon = 0.01$, and draw 100 paths simulated from the KMR Markov process for $N = 50$ players.

Figure 9 depicts the resulting mean sample paths. The evolution of the repricing economy follows an intuitive pattern. Initially, the undercutting strategy is the most attractive to agents: their main competitors are fixed-pricers which this strategy easily beats. However, a problem emerges as undercutters become more common: when they are matched against each other, they quickly compete down to costs and stay there. We soon see agents switching to cycling strategies to avoid this issue. These strategies allow agents to extract rent when competing with undercutters and fixed-price agents. However, the cycling strategy does not perform well against other cyclers. As discussed above, this is because of the cycling strategy's inability to punish an opponent that prices below the intended cycle. This inability ensures that we see the market converging to a mix of undercutting and cycling strategies.

Finally, the example satisfies the conditions of Claim 5 above: and indeed, in the right panel of Figure 9 the mean price dips briefly after initially increasing, only to shoot up eventually.

### 4.4 Empirics

The above discussion suggests that repricing may initially increase welfare; however, if undercutting turns into cycling, prices will increase, and welfare will decrease. So did cycling increase over time? We employ historical data from Keepa[17] to answer this question. The advantage of the Keepa data is that it extends back to 2011 (as opposed to our proprietary data that extends only to 2018).

The dataset we employ consists of the entire price history of the 10,000 best-selling products (as of July 2019). We use our price war and cycle recognition algorithms for each year and product separately to generate a dataset at the product-year level. We then run regressions of various measures of interest on product-fixed effects and year dummies, i.e.

$$y_{it} = \alpha_t + \beta_i + \epsilon_{it}$$

where $\alpha_t$ is a set of year dummies and $\beta_i$ is a set of product fixed effects. We report the coefficients on the year dummies in Figure 10. The dashed grey line marks the introduction of the Amazon MWS Subscription API, which allows repricers to subscribe to pricing change notifications. As we can see in the figure, the number of price wars and cycles have increased dramatically over time after MWS API's introduction.

[17]Keepa is a company providing API access to scraped data from Amazon Marketplace. We emphasize that this dataset differs from the data employed in the rest of the paper.

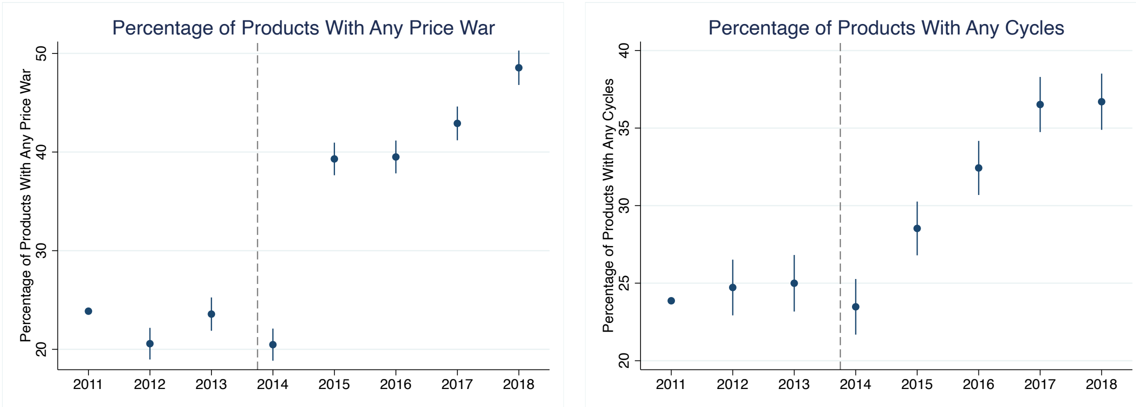


Figure 10: The Secular Increase in Price Wars and Cycling.
*Notes*: This figure illustrates how both cycling and price wars have increased since Amazon introduced the Subscription API in September 2013 (sellercentral.amazon.com). Each panel plots the coefficient and standard errors (clustered at the product level) from regressing a measure of cycling or price wars on product fixed-effects and year fixed-effects.

# 5 Conclusion

This paper has employed unique high-frequency e-commerce data, a novel algorithmic cycle-recognition approach, and a model of equilibrium in delegated strategies to argue that automated repricing on e-commerce platforms may have profound welfare implications. Firstly, we provided evidence that when a merchant starts repricing, he is likely to substantially lower his prices by essentially undercutting the lowest price in the market by the smallest possible amount. Secondly, to avoid the stark Bertrand-Nash competition of multiple undercutting merchants meeting, repricing companies have developed resetting strategies that regularly raise the price if it has fallen below some value. Thirdly, exploiting the fact that our source repricing company sometimes becomes overwhelmed with the number of scheduled repricing requests (and hence fails to execute some), we found evidence that these resetting strategies effectively raise opponents' prices. Finally, resets happen much more frequently during those hours of the night when sales probabilities are lowest, indicating that repricers are aware of the strategic tradeoffs inherent to their resetting strategies.

The resetting strategies create cycles reminiscent of Maskin-Tirole's Edgeworth cycles. However, while the theoretical literature has focussed on the possibility of Edgeworth cycles emerging as Markov-Perfect equilibria, we provide evidence that the cycles on Amazon are better understood as an equilibrium in *delegated* strategies. Thus, the chosen prices are not necessarily optimal in each period. Indeed, cycles are reset either when the price reaches a pre-specified level (with no war of attrition) or at a particular time of the day (particularly at night). Furthermore, the interfaces in which merchants have to enter their pricing rules allow *only* deterministic resets, effectively making it impossible to implement Maskin-Tirole strategies.

To what extent can delegated cycling strategies be sustained in equilibrium? To answer this question, we built a model of equilibrium in delegated strategies. We find that *miscoordination* between cyclers and undercutters is a critical ingredient in sustaining cycling. Nevertheless, the market can support a stable mix of these two strategies in equilibrium in an evolutionary model. Thus, the required miscoordination remains plausible as a long-run outcome. Furthermore, our model predicts that even though introducing repricing strategies may initially lower prices, it can dramatically increase prices in the long run. Crucially, unlike Maskin-Tirole, our model suggests prices are not anchored by the need for a war of attrition at marginal cost. Instead, firms choose the cycle such that the average price is close to a monopolist's price.

Our results have implications for managers selling products on online marketplaces and managers in charge of designing and policing such marketplaces. As a seller, our results suggest that adopting an appropriate repricing algorithm can potentially raise profits. In particular, undercutting algorithms successfully exploit the fact that competitors rarely update their prices. Similarly, resetting algorithms are effective when a competitor is known or suspected of employing an undercutting algorithm. As a platform designer, a manager needs to be aware of the draw that algorithmic repricing holds for merchants. Given its potentially deleterious impact on consumer welfare, policing frequent price changes could become necessary. Still, any gains from such policing must be traded off against the potential drawbacks from less efficient prices if marginal costs fluctuate over time.

## References


Abadie, A., S. Athey, G. Imbens, and J. Wooldridge (2017). When Should You Adjust Standard Errors for Clustering? *arXiv*.

Albaek, S., P. Mollgaard, and P. B. Overgaard (1997, December). Government-Assisted Oligopoly Coordination? A Concrete Case. *Journal of Industrial Economics 45*(4), 429–443.

Anonymous (2022). Algorithmic Pricing Facilitates Tacit Collusion: Evidence from E-Commerce. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, New York, NY, USA, pp. 32–33. Association for Computing Machinery.

Asker, J., C. Fershtman, and A. Pakes (2022). Artificial Intelligence, Algorithm Design, and Pricing. *AEA Papers and Proceedings 112*, 452–456.

Assad, S., R. Clark, D. Ershov, and L. Xu (2020). Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market.

Bank, E. C. (2004, April). How Frequently Do Prices Change? Evidence Based on the Micro Data Underlying the Belgian CPI. pp. 1–74.

Brown, Z. and A. McKay (forthcoming). Competition in Pricing Algorithms. *American Economic Journal: Microeconomics*.

Byrne, D. P. and N. d. Roos (2019, 02). Learning to Coordinate: A Study in Retail Gasoline. *American Economic Review 109*(2), 591 – 619.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review 110*(10), 3267–3297.

Chen, L., A. Mislove, and C. Wilson (2016). An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace. In *the 25th International Conference*, New York, New York, USA, pp. 1339–1349. ACM Press.

Eckert, A. (2003, February). Retail Price Cycles and the Presence of Small Firms. *International Journal of Industrial Organization 21*(2), 151–170.

Edelman, B. and M. Ostrovsky (2007). Strategic Bidder Behavior in Sponsored Search Auctions. *Decision Support Systems 43*(1), 192–198.

Ellison, G. and S. F. Ellison (2009). Search, Obfuscation, and Price Elasticities on the Internet. *Econometrica 77*(2), 427–452.

Green, E. J. and R. H. Porter (1984, January). Noncooperative Collusion under Imperfect Price Information. *Econometrica 52*(1), 87.

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

Harrington, J. E. (forthcoming). The Effect of Outsourcing Pricing Algorithms on Market Competition. *Management Science*.

Hausman, C. and D. S. Rapson (2018). Regression Discontinuity in Time: Considerations for Empirical Applications. *Annual Review of Resource Economics 10*(1), 533–552.

Johnson, J., A. Rhodes, and M. R. Wildenbeest (2021). Platform Design When Sellers Use Pricing Algorithms. *SSRN Electronic Journal*.

Kandori, M., G. J. Mailath, and R. Rob (1993, January). Learning, Mutation, and Long Run Equilibria in Games. *Econometrica 61*(1), 29–56.

Kandori, M. and R. Rob (1995, April). Evolution of Equilibria in the Long Run: A General Theory and Applications. *Journal of Economic Theory 65*(2), 383–414.

Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics 52*(3), 538–558.

Lee, K.-H. and L. Musolff (2021). Entry into Two-Sided Markets Shaped By Platform-Guided Search. *Working Paper*.

Leisten, M. (2022). Algorithm competition, with Humans. *Working Paper*.

Levy, D., M. Bergen, S. Dutta, and R. Venable (1997, August). The Magnitude of Menu Costs: Direct Evidence From Large U. S. Supermarket Chains. *The Quarterly Journal of Economics 112*(3), 791–825.

Maskin, E. and J. Tirole (1988). A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles. *Econometrica 56*(3), 571.

McCrary, J. (2008, 02). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics 142*(2), 698–714.

Miklós-Thal, J. and C. Tucker (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science 65*(4), 1552–1561.

Milgrom, P. and J. Roberts (1990, November). Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica 58*(6), 1255–1277.

Noel, M. D. (2007, March). Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market. *The Journal of Industrial Economics 55*(1), 69–92.

Noel, M. D. (2008, June). Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria. *Journal of Economics & Management Strategy 17*(2), 345–377.

Noel, M. D. (2018). Edgeworth Price Cycles. In *The New Palgrave Dictionary of Economics*, pp. 3463–3468. London: Palgrave Macmillan, London.

Popescu, D. (2015). Repricing Algorithms in E-Commerce. *SSRN Electronic Journal*.

Rotemberg, J. J. and G. Saloner (1986, June). A Supergame-Theoretic Model of Price Wars during Booms. *The American Economic Review 76*(3), 390–407.

Salcedo, B. (2015, 11). Pricing algorithm and collusion. pp. 1 – 30.

Schlosser, R. and M. Boissier (2017, September). Optimal Price Reaction Strategies in the Presence of Active and Passive Competitors. In *6th International Conference on Operations Research and Enterprise Systems*, pp. 47–56. SCITEPRESS - Science and Technology Publications.

Wang, Z. (2009a, December). (Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation. *Journal of Political Economy 117*(6), 987–1030.

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

Zhang, X. (2005). Finding Edgeworth Cycles in Online Advertising Auctions. *Mimeo*.

# A Cycling by Category

Figure A.1 shows that some product types are more likely to cycle than others. Furthermore, the product categories more likely to cycle generally match categories over-represented in our repricer-sourced product sample.

What factors influence the likelihood of cycling? Recall from Noel (2008) that cycling seems easier to maintain if there are high own-price elasticities and weak capacity constraints. This aligns well with our findings. In particular, while 'small' items such as books or DVDs (with negligible storage costs and capacity constraints) cycle often, appliances (which take up more space) cycle much less frequently.

However, one additional factor relative to Noel (2008) influences cycling propensity in our data: the extent to which the relevant competition for a product is other offers on the same platform. For instance, books are not very substitutable with each other. Therefore, the relevant competition for booksellers is other merchants making offers on the same book. Furthermore, customers expect Amazon to offer reasonable prices on books, and hence comparison shop only infrequently. Thus, the relevant competition is merchants making offers on the same book on the same platform. But it is precisely these merchants for which pricing according to strategies like those employed by the repricer is feasible: if the relevant competition included offers on a set of products on the same platform that were close substitutes, Amazon's API would not provide the merchants with access to the required monitoring technology to effectively implement their strategies[18]. This may provide a further reason why big-ticket items (like Electronics and Appliances) are less likely to cycle: customers are more likely to comparison-shop for these items, thereby ensuring that offers by merchants on other platforms – offers, i.e., for which Amazon's API provides no information[19] – must be considered amongst the leading competitors.

[18]Some repricers have started offering 'cross-product' repricing, but the technique seems not to have caught on, not least because under the API rules the merchant needs to at least pretend to be selling each product for which he wants to obtain regular pricing information – and being caught creating 'fake' offers can be extremely costly.

[19]As of April 2019, Amazon has started providing some information about what it currently considers to be competitive prices for a given product on other platforms through its API. Still, this data is naturally lower quality than the data on Amazon offers. However, it may be possible for repricers to maintain cross-platform cycles in the future.

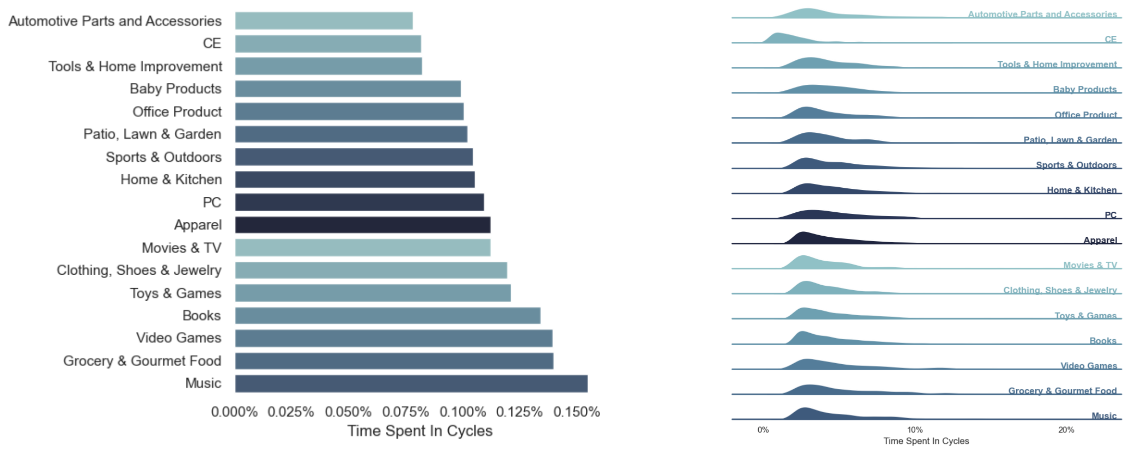


(a) Average (b) Distribution Conditional On Not Zero

Figure A.1: Fraction of Time Cycling By Product Category.

*Notes*: We compute the fraction of time spent in cycles (averaged across all offers for a product and across products). The left panel shows the mean fraction of time spent cycling, and the right panel shows the product-level distribution of time spent cycling, conditional on at least one cycle.

# B Proofs

**Proposition 1.** *Let* $(\ell^*, u^*) := \arg\max V_r(\ell, u)$ *and* $x^* := \min\{x : \pi(x) = V_r(\ell^*, u^*)\}$*. When a* $U$ *agent faces a* $C$ *agent, there is a continuum of equilibria given by*

$$\left\{\left((\underline{p}_U), (\underline{p}_C, \overline{p}_C)\right) : \underline{p}_U \leq x^*, \underline{p}_C = \ell^*, \overline{p}_C = u^*\right\}.$$

*These equilibria all result in identical price paths and payoffs.*

*Proof.*

1. Suppose $\underline{p}_C \geq \overline{p}_C$. Then $p_C \equiv \overline{p}_C$ for all values of $p_U$. This cannot be part of an equilibrium:

   (a) Suppose $\overline{p}_C > c$. Then, in equilibrium, $U$ must have $\underline{p}_U \in [0, \overline{p}_C)$. Thus, $C$ is making zero profits and would be better off by setting $\overline{p}_C$ arbitrarily high (ensuring a cycle with positive profits).

   (b) Suppose $\overline{p}_C = c$. Then, no matter what $\underline{p}_U$ is played, $C$ could deviate to playing $\overline{p}_C = 2\underline{p}_U$ for positive profits (he was previously making zero profits).

2. Suppose $\underline{p}_U > \underline{p}_C$. This cannot be part of equilibrium: either $C$ is making weakly negative profits (if $\underline{p}_C \leq c$) and could deviate to making positive profits by increasing $\underline{p}_C$, or $U$ (currently making zero profits) has space to set $\underline{p}'_U \in (c, \underline{p}_C]$ for positive profits.

3. We are left with $\overline{p}_C > \underline{p}_C \geq \underline{p}_U$. There is no profitable deviation for $U$: the (positive) payoff is the same as long as $\underline{p}_U \leq \underline{p}_C$ and if $\underline{p}_U > \underline{p}_C$ it becomes zero. To ensure that there is no profitable deviation for $C$, we need to make sure (a) that $C$ is choosing the optimal cycle and

(b) that $C$ does not want to capture the whole market at constant price $\underline{p}_U - \epsilon$ for arbitrarily small $\epsilon$. If $\underline{p}_U > x^*$, $C$ prefers to capture the whole market. So $\underline{p}_U \leq x^*$. When choosing to cycle, $C$ will choose $(\underline{p}_C, \overline{p}_C) = (\ell^*, u^*)$. Note that $\ell^* > x^*$ as $\ell < p_m$, $x^* < p_m$ and

$$\pi(\ell^*) = \ 2V_r(\ell^*, u^*) > V_r(\ell^*, u^*) = \pi(x^*)$$

where the first equality follows from the FOC for $\ell^*$. ☐

**Proposition 2.** *If $\pi(\cdot)$ is single-peaked, the solutions to* $\arg\max V_r(\ell, u)$ *must satisfy the FOCs.*

*Proof.* We verify that the Hessian is negative definite. Firstly, note

$$\frac{\partial V_r}{\partial u^2} = \frac{1}{2(u-\ell+k)}\Big\{\pi'(u) - 4\frac{\partial V_r}{\partial u}\Big\} < 0$$

as $\frac{\partial V_r}{\partial u} = 0$ and $\pi'(u) < 0$ are implied by the FOCs given the single-peakedness of $\pi$. Secondly, the determinant of the Hessian of $V_r(\cdot, \cdot)$ satisfies

$$\begin{aligned} D = -\frac{1}{4(u-\ell+k)^4}\Big\{ \quad & (\pi(\ell) - \pi(u))^2 \\ & -4(u-\ell+k)^2\Big(\pi'(u)\frac{\partial V_r}{\partial \ell} + \pi'(\ell)\frac{\partial V_r}{\partial u}\Big) \\ & +(u-\ell+k)^2\pi'(\ell)\pi'(u) \quad \Big\}, \end{aligned}$$

At the FOCs, the term on the second line evaluates to zero and further $\pi(\ell) = \pi(u)$ so that we are left with only the last line. As $\pi$ is single-peaked, $\ell^* < p_m < u^*$ and hence $D > 0$. ☐

**Proposition 3.** *If $\pi$ is symmetric around the monopoly price $p_m$ and $\tilde{\pi} = g \circ \pi$ where $g'' \geq 0, g' > 0$ and $g(0) = 0$, then $\tilde{\ell}^* \geq \ell^*$ and $\tilde{u}^* \leq u^*$.*

*Proof.* By Jensen's inequality

$$\begin{aligned} \frac{1}{u-\ell+k}\int_\ell^u g\Big(\pi(p)\Big)dp \ &\geq \ \frac{u-\ell}{u-\ell+k} g\Big(\frac{1}{u-\ell}\int_\ell^u \pi(p)dp\Big) \\ &\geq \ g\Big(\frac{1}{u-\ell+k}\int_\ell^u \pi(p)dp\Big) \\ &= \ g(\pi(u)) = g(\pi(\ell)). \end{aligned}$$

Thus either $\ell$ must increase or $u$ decrease, and as $\tilde{\pi}(\cdot)$ will also be symmetric around the same $p_m$ we thus have that both $\ell$ must increase and $u$ decrease. ☐

**Proposition 4.** *If $\pi(\cdot)$ is symmetric around the monopoly price $p_m$, the average price during the decreasing phase of the cycle will satisfy $p_{avg} = p_m$.*

*Proof.* From the FOCs, $\pi(u) = \pi(\ell)$. The symmetry of $\pi$ then implies

$$p_m - \ell = u - p_m \implies p_m = \frac{u+\ell}{2}.$$

□

**Proposition 5.** *Suppose that in each period only one player is allowed to adjust their strategy. Then the limit distribution puts probability 1/2 on* $(0, \lfloor \gamma N \rfloor, N - \lfloor \gamma N \rfloor)$ *and probability 1/2 on* $(0, \lceil \gamma N \rceil, N - \lceil \gamma N \rceil)$.

*Proof.* In the notation of KMR: as $F$ is never a best response, we have $v_z \geq 1$ for any $z = (n_F, n_U, n_C)$ such that $n_F > 0$. However, as in KMR, $z = (0, \gamma N, (1-\gamma)N)$ is a global attractor for any $b(\cdot)$ and thus $v_{\gamma N} = 0$. □

**Corollary 1.** *The mean price under the limit distribution is given by*

$$p_\infty = [1 - 2\gamma(1-\gamma)]c + 2\gamma(1-\gamma)p_m,$$

*where* $p_m$ *is the monopoly price and* $c$ *is the unit cost.*

*Proof.* This follows trivially as in equilibrium exactly fraction $\gamma$ of repricers are playing U and $1-\gamma$ are playing C. The price is $c$ whenever U and U or C and C meet, and it is $p_m$ (on average) when $C$ and $U$ meet. □

**Proposition 6.** *If* $p_\infty > p_0$ *and* $c_C - c_U$ *is sufficiently high relative to the dispersion of the fixed pricers' price distribution* $F(\cdot)$*, then prices will first increase, then decrease and then increase again. Furthermore, prices will dip below the initial price.*

*Proof.* Suppose $\alpha_i$ is the fraction of the population playing strategy $i$. As long as

$$\alpha_C V_{nr}^* - c_U > \alpha_U V_r^* - c_C,$$

the best-response to the population distribution of strategies will be U. Note that initially, $\alpha_F = 1$ and hence this inequality must be satisfied as $\epsilon \to 0$. Furthermore, the inequality will keep being satisfied as long as

$$\alpha_U < \alpha_U^* \equiv \frac{c_C - c_U}{V_r^*}.$$

But when the economy hits this boundary, the mean price will be given by

$$\mathbb{E}[p_t] = (\alpha_U^*)^2 \times 0 + 2\alpha_U^*(1-\alpha_U^*)\mathbb{E}[\tilde{p}] + (1-\alpha_U^*)^2\mathbb{E}[\min\{\tilde{p}_1, \tilde{p}_2\}].$$

This will be lower than the starting price as long as

$$\frac{2\alpha_U^* - (\alpha_U^*)^2}{2\alpha_U^* - 2(\alpha_U^*)^2} > \frac{\mathbb{E}[\tilde{p}]}{\mathbb{E}[\min\{\tilde{p}_1, \tilde{p}_2\}]},$$

which will be the case if the costs of the repricing types are sufficiently different relative to the size of the support of the fixed pricers prices. □

## C Omitted Stage Game Equilibria

**Proposition 7.** *The unique equilibrium amongst undercutting agents is given by* $\underline{p}_1 = \underline{p}_2 = c$.

*Proof.* As we have

$$u_i(\underline{p}_i; \underline{p}_j) = \begin{cases} \pi(\underline{p}_j) & \text{if } \underline{p}_i < \underline{p}_j, \\ 0.5\pi(\underline{p}_j) & \text{if } \underline{p}_i = \underline{p}_j, \\ 0 & \underline{p}_i > \underline{p}_j, \end{cases}$$

the BR correspondence is is given by

$$BR_i(\underline{p}_j) = \begin{cases} [0, \underline{p}_j) & \text{if } \underline{p}_j > c \\ \mathbb{R}^+ & \text{if } \underline{p}_j = c. \end{cases}$$

□

**Proposition 8.** *If $k$ is sufficiently small, the unique equilibrium amongst two resetting agents is given by*

$$\Big((\overline{p}_1, \underline{p}_1), (\overline{p}_2, \underline{p}_2)\Big) = \Big((c, c), (c, c)\Big).$$

*Proof.* We proceed by ruling out all other equilibria.

1. If $\underline{p}_i \geq \overline{p}_i > c$, then $p_i \equiv \overline{p}_i$ whence a profitable deviation is to set $\underline{p}_j = \overline{p}_j = \overline{p}_i - \epsilon$ for some small $\epsilon$. Hence, this cannot be part of an equilibrium.

2. If $\underline{p}_i > \overline{p}_j > c$, then $p_j \equiv \overline{p}_j$ whence a profitable deviation is to set $\underline{p}_i = \overline{p}_i = \overline{p}_j - \epsilon$ for some small $\epsilon$. Hence, this cannot be part of an equilibrium.

3. If $\overline{p}_i > \overline{p}_j > \underline{p}_i > \underline{p}_j$, then note that player $i$ is choosing the lower end of the cycle and his choice $\underline{p}_i$ must hence satisfy $\frac{\partial V_r(\underline{p}_i, \overline{p}_j)}{\partial \underline{p}_i} = 0$, i.e.

$$\pi(\underline{p}_i) = \frac{1}{\overline{p}_j - \underline{p}_i + k} \int_{\underline{p}_i}^{\overline{p}_j} \pi(p) dp. \tag{4}$$

   To ensure that player $j$ does not have a profitable deviation to $\overline{p}_j = \underline{p}_j = \underline{p}_i - \epsilon$, we must have that $V_{nr}(\overline{p}_j, \underline{p}_i) \geq \pi(\underline{p}_i)$. Using (4), this simplifies to

$$V_{nr}(\overline{p}_j, \underline{p}_i) \geq 2V_r(\overline{p}_j, \underline{p}_i).$$

   As $\lim_{k \to 0} V_{nr}(\overline{p}_j, \underline{p}_i) = V_r(\overline{p}_j, \underline{p}_i)$ this must be violated for sufficiently small $k$.

4. If $\overline{p}_i > \overline{p}_j > \underline{p}_j > \underline{p}_i$, then note that player $j$ is resetting the cycle and also choosing the upper end of the cycle. This can never be part of equilibrium as the resetting party has a strictly

higher incentive to extend the cycle upwards:

$$\frac{\partial V_r}{\partial u} - \frac{\partial V_{nr}}{\partial u} = \frac{k}{(u-\ell+k)^2}\pi(u) - \frac{k}{u-\ell+k}\pi'(u) > 0.$$

□

# D Cycle & Price War Recognition

## D.1 Cycle Recognition

This paper differs from past literature on cycling first and foremost in the size of the data available to us. While previous studies have typically conducted targeted data collection in just a few markets, our approach employs a large quantity of data on prices across hundreds of thousands of markets on an e-commerce platform. To make things worse, due to the low setup costs involved, competition in e-commerce is subject to frequent entry and exit – phenomena that can suddenly interrupt cycling behavior.

Given these complications, we require a machine-legible cycle definition to analyze cycling in this data. Any attempt at hand-coding cycling behavior is doomed to failure. Choosing a fixed definition of cycling makes the classification problem feasible and has the further benefit of disciplining the analysis by forcing us to use consistent criteria across markets. As we have the luxury of a large amount of data, we choose to minimize the number of false positives by providing a rather strict definition of cycling behavior.

**Definition 1.** *A sequence of prices between two local maxima is deemed a cycle if and only if*

1. *the price is strictly decreasing between the local maxima,*
2. *the reaction time never exceeds 10x the average,*
3. *the distance between maxima is at most 40% of the amplitude,*
4. *the distance between minima is unconstrained, and*
5. *there are at least 7 steps in the cycle.*

*These criteria will be relaxed by 2% with each consecutive cycle, but the relaxation will top out when the criterion is multiplied by 5. Furthermore, we will only keep cycles that are part of a run of at least four cycles.*

To illustrate which type of behavior this definition captures, we provide an example in Figure D.1. The blue line indicates the landed price for a given offer, and a green backdrop highlights periods the algorithm identifies as cycling. We can see that while the offer is identified as cycling most of the time, there are some exceptions: e.g., at 9 AM on August 29, the offer drops almost all the size of the cycle in just one step. The classification algorithm expects each cycle to have at least seven steps and hence considers this anomaly indicative of an end to cycling. More realistically,

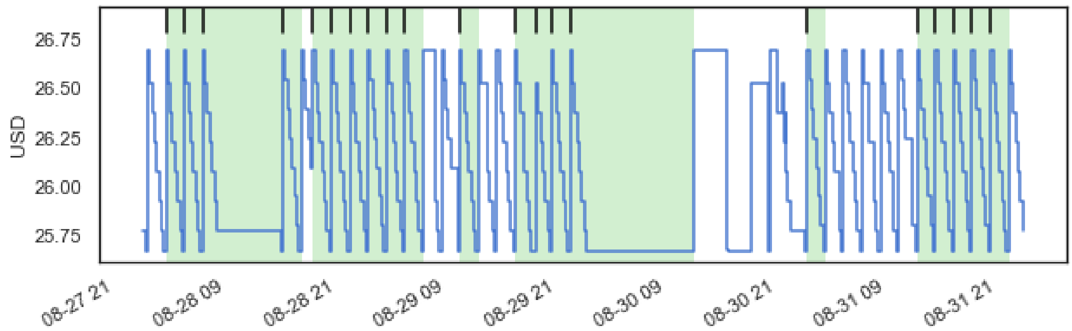


Figure D.1: Output of The Cycle Recognition Algorithm.
*Notes*: This figure illustrates the output of the cycle recognition algorithm using an example price path. The periods shaded in green are declared to be ones in which the offer is cycling by the algorithm.

Amazon's servers were busy and delayed in notifying the repricer of the usual regular price changes. Less obviously, the same problem occurs on August 31 (the cycle briefly takes six rather than seven steps). As we can see, the algorithm's cautious approach leads to many false negatives but also heavily restricts the rate of false positives. Furthermore, as we can see, e.g., on August 30 in the Figure, the algorithm does not consider comparatively long stops at the bottom of the cycle evidence against cycling: this is important because we want to avoid ruling out a war of attrition ex-ante.

Having classified individual offers as cycling or not cycling during a specific period, we can exploit the fact that cycling is competitive behavior to validate and refine our classification algorithm. According to theory, we should never observe merchants cycling independently: if there is any cycling, there should always be at least two cycling merchants. This fact helps us verify to what extent the classification algorithm is misled by noise: when fooled in this way, the probability of observing a cycle in one offer would be conditionally independent of that of observing a cycle in another offer on the same product. Reassuringly, this independence hypothesis is soundly rejected by our data: we find that there is precisely one offer that is cycling 2.58% of the time. However, conditional on at least one cycling offer, the probability that there are two cycling offers is 26.38%. Keeping in mind that we do not observe all offers (only the top twenty), these figures suggest that the algorithm picks up mostly actual cycling even though it is inevitably picking up a fair amount of noise along the way.

## D.2 Price War Recognition

We utilize the following definition of a price war:

**Definition 2.** *A sequence of prices between is deemed a price war if and only if*

1. *the price is strictly decreasing for at least 7 consecutive steps,*
2. *the price declines at least \$0.01 and at most \$5.00 at each step,*
3. *the time between subsequent steps never exceeds 10h.*

| Dep. Var. \ Bandwidth | ±19 | ±15 | ±10 | ±5 |
|---|---|---|---|---|
| Log Offer Price | -0.041** | -0.043*** | -0.041** | -0.052*** |
| | (0.016) | (0.014) | (0.016) | (0.012) |
| Log Buybox Price | -0.000 | 0.000 | -0.001 | -0.007*** |
| | (0.003) | (0.002) | (0.002) | (0.002) |
| Buybox Share | 0.040** | 0.042** | 0.049** | 0.049*** |
| | (0.017) | (0.017) | (0.019) | (0.015) |
| Log(Profits + 1) | 0.325*** | 0.425*** | 0.482*** | 0.426*** |
| | (0.095) | (0.113) | (0.110) | (0.127) |

Table E.1: RD Estimates for Balanced Panel.

*Notes*: This table provides the $\hat{\beta}$ of (1), i.e., the estimated effect of repricer activation on various outcome variables (row) and for various estimation windows (column). The difference to Table 1 is that here we subset to a balanced panel. For instance, in the column "±19", we now use only data from offers which we observed for the full 38 day window around their repricer activation date. All standard errors are clustered at the merchant level. *** p<0.01, ** p<0.05, * p<0.1

# E Robustness of RD

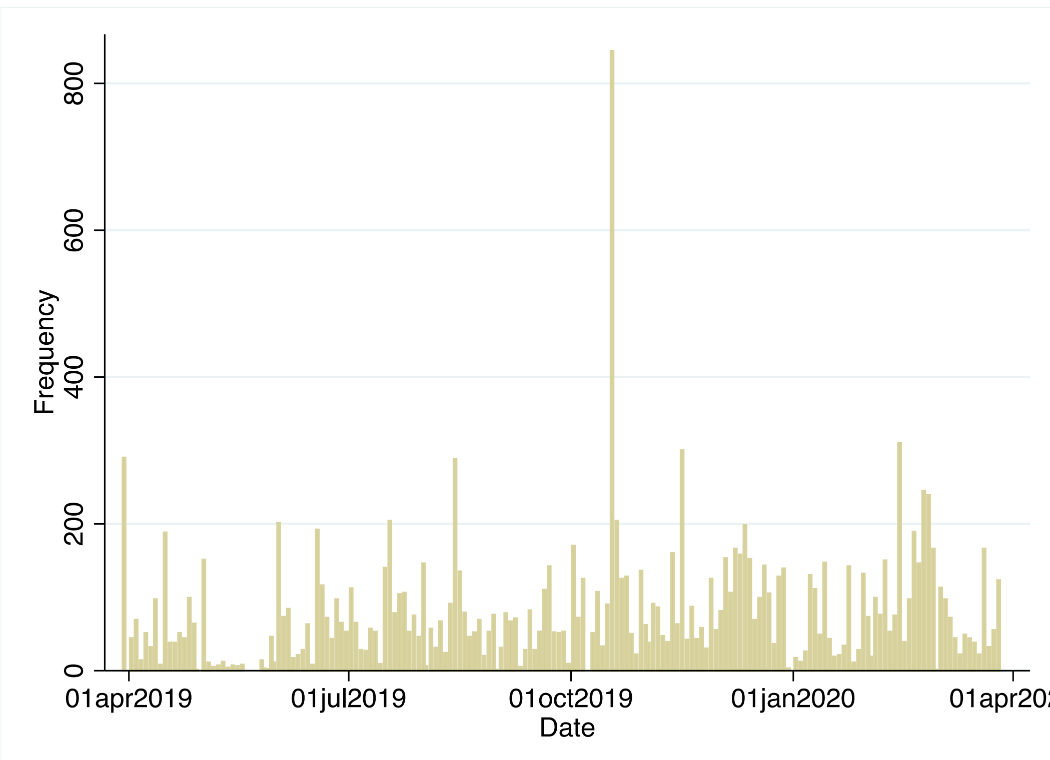


Figure E.1: Distribution of Repricer Activation Events

*Notes*: This figure illustrates distribution of repricer activation events used in the estimation of the repricer activation RD.

| Dep. Var. \ Bandwidth | ±19 | ±15 | ±10 | ±5 |
|---|---|---|---|---|
| Log Offer Price | -0.010*** | -0.011*** | -0.011*** | -0.019*** |
| | (0.002) | (0.002) | (0.004) | (0.007) |
| Log Buybox Price | -0.002*** | -0.003** | -0.004** | -0.006*** |
| | (0.001) | (0.001) | (0.001) | (0.002) |
| Buybox Share | 0.016*** | 0.016*** | 0.014** | 0.025*** |
| | (0.004) | (0.004) | (0.006) | (0.008) |
| Log(Profits + 1) | 0.294*** | 0.290*** | 0.268*** | 0.224* |
| | (0.056) | (0.060) | (0.063) | (0.121) |

Table E.2: RD Estimates Controlling for Lagged Dependent Variable.

*Notes*: This table provides the $\hat{\beta}$ of (1), i.e., the estimated effect of repricer activation on various outcome variables (row) and for various estimation windows (column). The difference to Table 1 is that here we add a lag of the dependent variable as a control to each regression. All standard errors are clustered at the merchant level. *** p<0.01, ** p<0.05, * p<0.1

| Dep. Var. \ Bandwidth | ±19 | ±15 | ±10 | ±5 |
|---|---|---|---|---|
| Log Offer Price | -0.052*** | -0.050*** | -0.053*** | -0.009 |
| | (0.010) | (0.011) | (0.008) | (0.031) |
| Log Buybox Price | -0.006 | -0.009 | -0.016** | -0.012 |
| | (0.005) | (0.005) | (0.006) | (0.011) |
| Buybox Share | 0.041*** | 0.034*** | 0.024* | -0.016 |
| | (0.010) | (0.013) | (0.013) | (0.032) |
| Log(Profits + 1) | 0.243*** | 0.230*** | 0.134 | -0.844 |
| | (0.074) | (0.085) | (0.147) | (0.611) |

Table E.3: RD Estimates from a Quadratic Specification.

*Notes*: This table provides the $\hat{\beta}$ of (1), i.e., the estimated effect of repricer activation on various outcome variables (row) and for various estimation windows (column). The difference to Table 1 is that here we allow the dependent variable to depend *quadratically* on the running variable (still with separate coefficients to the left and right of the cutoff). All standard errors are clustered at the merchant level. *** p<0.01, ** p<0.05, * p<0.1

| Dep. Var. \ Bandwidth | ±19 | ±15 | ±10 | ±5 |
|---|---|---|---|---|
| Log Offer Price | -0.005*** | -0.001 | 0.004* | 0.003 |
| | (0.002) | (0.002) | (0.002) | (0.002) |
| Log Buybox Price | -0.003*** | -0.002** | -0.001 | -0.001 |
| | (0.001) | (0.001) | (0.001) | (0.001) |
| Buybox Share | 0.006** | 0.002 | -0.001 | -0.002 |
| | (0.002) | (0.003) | (0.004) | (0.003) |
| Log(Profits + 1) | -0.012 | -0.043 | -0.075* | -0.057 |
| | (0.026) | (0.029) | (0.038) | (0.038) |

Table E.4: RD Estimates from a Placebo Treatment.
*Notes*: This table provides the estimated effect of a placebo treatment (i.e., all of the estimated effects should be zero). In particular, we randomly assign each offer a new placebo treatment date that is 5 days earlier or later than its actual treatment date (both with 50% probability). We then perform the same estimation procedure as in Table 1. All standard errors are clustered at the merchant level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

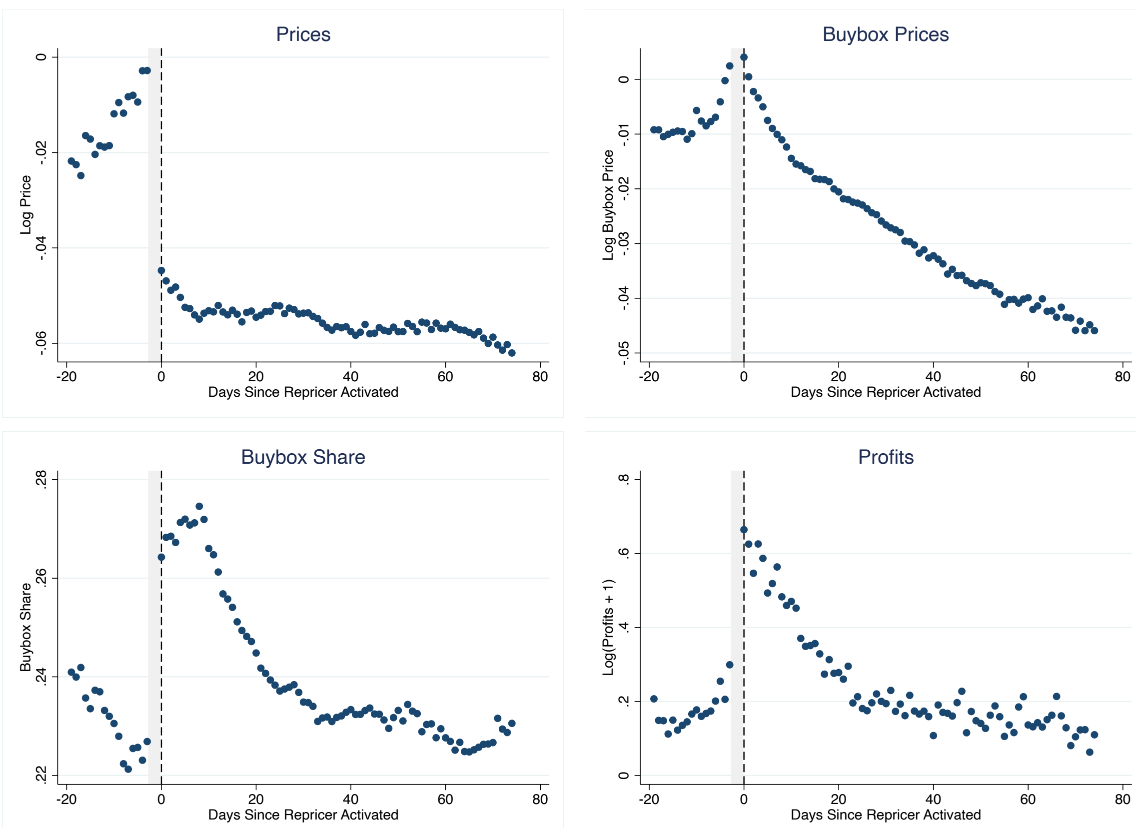


Figure E.2: Long-Term Impact of Repricing.
*Notes*: This figure extends the RD results from above to show the long-term trends that follow the activation of the repricer. The light-gray area indicates time during which treatment status cannot be ascertained due to sampling frequency.

# F Robustness of Resetting Queue

| | Dependent Variable | Reset Happened? | | $N$ | $R^2$ | Offer FE |
|---|---|---|---|---|---|---|
| (1) | Competitor # Price Changes | 0.399*** | (0.115) | 100,926 | 0.432 | Yes |
| (2) | Competitor Log Price | 0.0210*** | (0.00538) | 100,926 | 0.989 | Yes |

Table F.1: Price Resets Are More Effectice At Raising Serious Competitors' Prices.
*Notes*: This table shows the results of regressing various dependent variables (one per row) on an indicator for whether a price reset happened in the preceding day, using a sample of 858 products for which a price reset was configured but not always executed. In contrast to Table 2, this table only considers offers to be competitors if their average price (across the whole sample period) lies at most 5% above that of the resetting merchant. Our estimates show that these 'serious competitors' are much more likely to react to resetting by raising prices. Standard errors are clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1

| | Dependent Variable | Reset Configured? | | Reset Happened? | | Offer FE |
|---|---|---|---|---|---|---|
| (1) | Own # Price Changes | 0.356 | (0.312) | 1.597*** | (0.288) | Yes |
| (2) | Competitor # Price Changes | -0.119 | (0.102) | 0.275*** | (0.080) | Yes |
| (3) | Own Log Price | -0.037*** | (0.011) | 0.078*** | (0.010) | Yes |
| (4) | Competitor Log Price | -0.028*** | (0.006) | 0.011*** | (0.004) | Yes |
| (5) | Buybox Log Price | -0.038*** | (0.011) | 0.021*** | (0.007) | Yes |

Table F.2: The Effect of *Configuring* Price Resets Versus Price Resets Happening.

*Notes*: This table shows the results of regressing various dependent variables (one per row) on (i) an indicator for whether a reset was configured and (ii) an indicator for whether a reset happened. Thus, instead of conditioning on products with configured resets (like Table 2), this table controls for the potential selection problem by including an additional regressor. In doing so, we find confirmation for our hypothesis that the selection effect is negative, i.e., a reset is more likely to be configured on products where competitors are pricing more aggressively. Standard errors are clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1

| | (1) OLS | (2) 2SLS |
|---|---|---|
| Repricing Merchant Log Price | 0.213*** | 0.162*** |
| | (0.0183) | (0.0420) |
| First-Stage | | |
| Reset Happened? | | 0.0823*** |
| | | (0.0102) |
| First-Stage F-Value | | 65.00 |
| Observations | 156,853 | 156,853 |
| Mean Dependent Variable | 3.163 | 3.163 |
| Std. Dev. Dependent Variable | 0.748 | 0.748 |

Table F.3: There Is Significant Pass-Through of Resets to Competitors.

*Notes*: This table shows the result of regressing log competitor prices on log prices of the repricing merchant. The first column provides results from an OLS regression that might be contaminated by, e.g., cost shocks common to the repricing merchant and his competitors; hence, the coefficient should not be interpreted as causal. However, in the second column, we exploit the variation in whether the repricing company executes a reset. As execution fails for purely technological reasons, we argue in the main text that it satisfies an exclusion restriction, i.e., it should only affect competitor prices via repricing merchants' prices. Hence, the coefficient in the second column has a causal interpretation: about a sixth of the price increase due to resetting is passed on to competitor prices. Standard errors are clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1