Algorithmic Pricing, Price Wars and Tacit Collusion: Evidence from E-Commerce

Leon Musolff*

Mar 27, 2024

Abstract

As the economy digitizes, menu costs fall, and firms can more easily monitor prices. These trends have led to the rise of automated pricing (and re-pricing) tools. We employ a novel e-commerce dataset to examine the effect of algorithmic pricing in the wild. Evidence from an event study suggests that firms that start employing repricing tools drop their prices by 16.7%, with market prices falling by 9.5%. However, algorithmic pricing companies have developed 'resetting' strategies (which regularly raise prices in the hope that competitors will follow) in order to avoid stark Bertrand-Nash competition. We find that these strategies are effective at coaxing competitors to raise their prices: when a resetting strategy is adopted, both competitor prices and market prices eventually increase by 8%. While the resulting patterns of cycling prices are reminiscent of Maskin-Tirole's Edgeworth cycles, a model of equilibrium in delegated strategies fits the data better. This model suggests that the average price over the cycle will be the monopoly price. Moreover, if the available repricing technologies remain fixed, cycling and prices could rise significantly. However, cycling is still relative rare in the data.

JEL Classification: L41, L81, D47

*Wharton School of The University Of Pennsylvania (e-mail: lmusolff@wharton.upenn.edu). I am especially grateful to Jakub Kastl for many meetings full of insightful comments and stimulating discussions that helped shape this paper. I also thank Bruno Baranek, Stefano Baratuche, Nick Buchholz, Ana Luiza Dutra, Kate Ho, Kwok-Hao Lee, Adam Kapor, Nikita Melnikov, Jonas Mueller-Gastell, Christopher Neilson, Franz Ostrizek, Pellumb Reshidi. This paper benefited from comments by seminar participants at Princeton University, Simon Business School, Dartmouth, the Federal Trade Commission, National University of Singapore and anonymous referees as well as conference participants at both the Sapporo Workshop on Industrial Economics and ACM EC'22. Any errors are exclusively mine.

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 an extensive theoretical literature on dynamic pricing games, that body of work has shown that many price patterns are consistent with equilibrium conduct, as captured by the Folk Theorem. Hence, is it unclear what the price effects will be from a rise in the frequency with which firms can adjust their prices. In particular, will it intensify competition through price undercutting or facilitate sustaining high prices? Or will other patterns emerge? In light of the absence of precise predictions, empirical research is needed to inform us about the state of competition in online markets.

This paper employs data on the pricing decisions made by third-party sellers on Amazon Marketplace and a new model of equilibrium in delegated strategies to empirically assess and expand the theoretical predictions of the literature on dynamic pricing. We find that delegation of pricing to simple algorithms can lead to tough price wars but also facilitate tacit collusion. Furthermore, the algorithms currently employed are related to but distinct from 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 or strategy innovation, the market would fully transition from static to dynamic pricing. Along the transition path, prices would initially fall but eventually rise as the (average) equilibrium price between (some) repricing algorithms is the monopoly price. However, the simplicity of the algorithms is also their greatest weakness in that collusive profits crucially depend on a mix of repricing technologies being employed in the market: none of the observed algorithms performs well when playing against itself. Furthermore, potential tacit collusion could be interrupted by the entry of additional firms (which may be attracted to enter by the high prices). In summary, the analysis is able to characterize the rich pricing dynamics observed on Amazon Marketplace while identifying some of the key driving forces at play.

This paper proceeds as follows. In Section 1, we discuss prior literature. In Section 2, we provide details on the institutional setting (Amazon Marketplace), which is characterized by high own-price elasticities and ease of price adjustment. We also provide some initial evidence of interesting pricing patterns. In Section 3, we distill our understanding of repricing using evidence from two event studies. The first event study establishes that turning on repricing lowers a merchant's prices by 16.7%, initially without adversely affecting market prices, but in the long-term with a strong negative effect on market prices which eventually decline by 9.5%. The second event study zooms in on a particular kind of (currently still relatively rare) strategy that regularly raises a merchant's prices. We show that this strategy effectively coaxes competitors to raise their prices in turn, eventually increasing market prices by 8%. In Section 4, 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 Literature Review

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 post 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 elasticities can reach -20 (Ellison and Ellison, 2009), the prospect of low prices under competition could induce firms to collude.

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) just like humans (Byrne and Roos, 2019), at least under some learning protocols (Asker et al., 2022; Banchio and Mantegazza, 2022). Supra-competitive prices may be especially likely when algorithms have a misspecified model that does not directly account for the effect of competitors' prices on demand (Hansen et al., 2021; Cooper et al., 2015).

Still, these warnings are speculative and are largely based on simulation studies or theoretical analyses. By contrast, this paper contributes an empirical study of pricing 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, 2022). Indeed, gasoline prices in Germany increase only when *both* members of a local duopoly adopt algorithmic pricing strategies (Assad et al., 2024). Similarly, the algorithms 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 (with one algorithm reacting much faster to opponent price changes than the other) as this can effectively turn the game from Bertrand to Stackelberg (Brown and MacKay, 2023). 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 prices up. By contrast, managerial override – potentially a key feature of gasoline markets (Leisten, 2022) – is unlikely in our setting.

As we discuss below, our pricing data exhibits cycling patterns, which are reationalized by two kinds of models in the theoretical literature. Firstly, some consumer search models predict cycling. For instance, persistent marginal cost fluctuations can lead to price cycles under noncooperative play (Tappata, 2009). High costs engender less price dispersion, prompting consumers to search less. Less intensive search, in turn, lowers demand elasticity and raises cost pass-through. However, marginal cost fluctuations at the speed of the cycles we observe (mostly faster than daily, sometimes significantly) seem implausible. Similarly, cycles that occur if consumers strategically delay search (Fershtman and Fishman, 1992) require consumers to be able to forecast price movements. The ability to make such forecasts is implausible in our setting: there are too many products, their prices are too low to make it worthwhile to keep track of price movements, and the cycles are too fast.

Secondly, 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). Though MPE are a natural benchmark when humans make pricing decisions, this ceases to be realistic when prices change as often as in our dataset. Instead, merchants delegate short-term pricing decisions to a computer (Chen et al., 2016).

Thus, while delegating strategic decisions to agents facing a distorted payoff function can already achieve commitment (Fershtman et al., 1991), commitment in the algorithmic pricing world does not involve a strategic delegate. Instead, the delegation is to an algorithm which faithfully executes the principal's instructions. By contrast to the prior literature, the source of commitment is thus not a cleverly chosen objective function but rather a restriction in strategy space necessitated by the cost of communicating a strategy to the computer.

The limited vocabulary of the available repricing interfaces makes a crucial difference. This difference is explored 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 á 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 the monopoly price.

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

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

(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.

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[2].

The fact that there exist multiple offers for most products often goes ignored by customers: about 83% of purchases go via 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[3] (Lee and Musolff, 2023). Indeed, as shown in Figure 2, the Buybox algorithm is very price sensitive: being priced just 1% more than the cheapest offer lowers the probability of winning the Buybox by more than half.

While price is key in deciding Buybox ownership, other factors such as fulfillment method, seller rating, and shipping time also play a role. In particular, Lee and Musolff (2023) find that after switching to Amazon's fulfillment network (FBA), a seller could raise their price by 12.36% and still command the same Buybox share. Similarly, after moving from the 1st to the 99th percentile on dispatch time, a seller would have to lower their price by 9.91% to command the same Buybox share. Furthermore, Amazon offers have a 5.70% price advantage, and moving from the 1st to the 99th percentile on the amount (quality) of feedback yields a 1.28% (0.73%) price advantage.

From a seller's perspective, offers being pooled on a product page means that winning the Buybox is crucial. Moreover, as sellers know that winning is mostly a matter of having the lowest price (especially once other offer characteristics have been fixed), we might expect intense competition and prices close to marginal cost. Indeed, we can think of the Buybox as approximately emulating

[2]UPC stands for Universal Product Code; these codes are administered by the non-profit GS1.

[3]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., 2023).

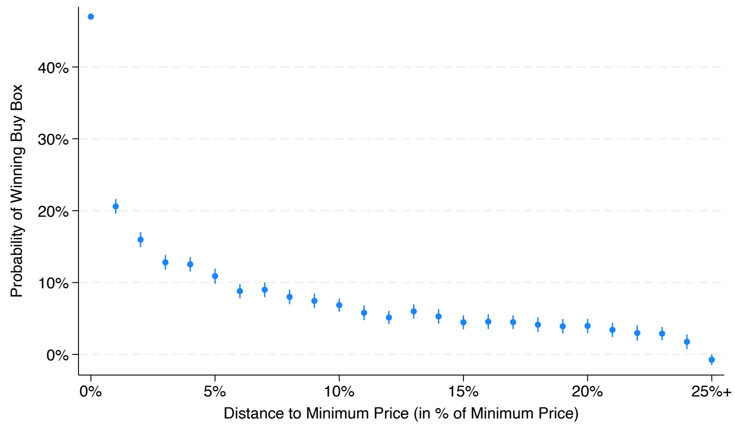

Figure 2: The Buybox is Very Price Sensitive.

Notes: This figure provides the probability of winning the Buybox if an offer is priced at a given percentage above a product's minimum priced offer; thus, a value of 0% on the horizontal axis corresponds to being (equally priced as) the cheapest offer. The probability of winning is very sensitive to price: being priced just 1% more than the cheapest offer lowers the probability by more than half. Formally speaking, the figure reports coefficients from a regression of an indicator for whether an offer won the Buybox on evenly-spaced bins of the percentage difference between the offer's price and the minimum priced offer's price. The regression also includes offer and market fixed effects so that the coefficient is identified only off variation that comes from an offer changing its relative price over time. Standard errors are clustered at the product level.

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 (relative to seller-reported costs and fees)[4] is 18.42%. Furthermore, there is considerable mass on even higher markups.

How reliable are our estimates of markup? The 18.42% median markup is consistent with a 18% median markup reported by 3,500 Amazon sellers in a survey[5] conducted by Jungle Scout in 2022 (Jungle Scout, 2023). However, both the survey and our data rely on seller-reported costs. To the extent that sellers report costs to the repricing company in an attempt to ensure there are no sales at a price below marginal costs, we would expect reported costs to include seller estimates of all costs incurred, including, e.g., the cost of labor, advertising, and storage. However, to guard against sellers not including other Amazon fees in their calculation (such as storage costs), we can replace the measured fees with a blanket 50% rate of fees paid to Amazon; we choose 50% as a conservatively high estimate because industry sources suggest a rate of 40%-45% for the period considered by our study (marketplacepulse.com). If we make this adjustment, the median markup falls to 10.63%. Alternatively, one may worry that about 18% of online purchases were returned in 2020 (cnbc.com).

[4]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 853,335 sales 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.)

[5]Yu (2023) also uses the 18% figure to calibrate his pricing model.

As Amazon does not refund fees for such returns, the correct markup should multiply reported fees by (1/0.82). With this adjustment, the median markup falls to 9.36%.

While our markups are an approximation which should be interpreted with caution, it is clear that they are far from zero. 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 frequently reprice 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 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 rules available 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 spent 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, one would expect prices to change frequently. 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.[6]

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 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 3, which depicts a typical Edgeworth-like cycle in our dataset. The product in question is a toddler t-shirt, and we depict the only three offers for this product between 2018-09-13 and 2018-10-30. The figure shows the offer's price paths and current possession of the Buybox (in the color strip at the bottom of each graph). Focusing on the bottom panel, which zooms in on the gray region 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.

However, once we examine the price paths in more detail, some discrepancies with the Maskin-

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

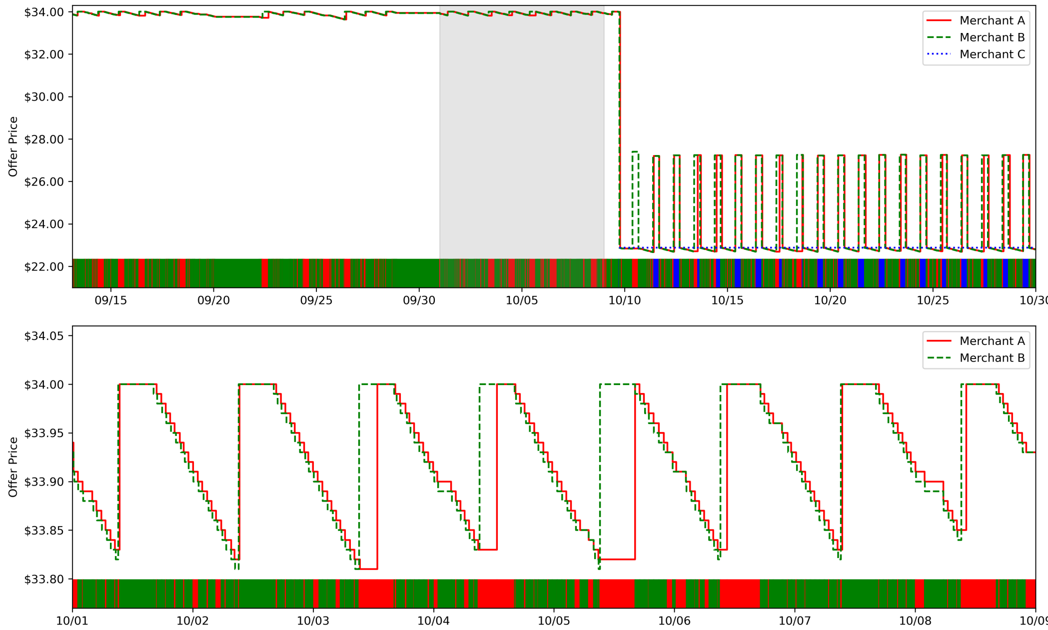

Figure 3: A Typical Cycle in the Pricing Data.

Notes: This figure depicts the prices of all offers for a toddler t-shirt between 2018-9-13 and 2018-10-30. The top panel shows the prices for the entire period, and the bottom panel zooms in on the period (indicated in gray in the top panel) between 2018-10-1 and 2018-10-9, just before the third merchant enters. During this period, both merchants' prices decrease slowly before jumping back up. During the price decrease, we see the characteristic 'undercutting' pattern where A regularly chooses a price just one cent below B's price. The bottom of each graph shows which offer has ownership of the Buybox at any given time.

Tirole theory begin to emerge. Merchant A has an advantage: often when the two sellers have the same price, A is allocated the Buybox. On the other hand, B must undercut A to gain the Buybox. Furthermore, A seems aware of this as it always chooses to match B's price while B is stuck strictly undercutting A'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, B immediately increases its price, and A follows, sometimes quickly and sometimes with some delay.

The asymmetry in pricing and Buybox dominance criteria suggests that there might be meaningful differences between the sellers. There is a reason why Amazon prefers to assign the Buybox to A when prices are tied: while both sellers' offers are fulfilled by Amazon, A has about double the number of ratings (622 vs 342) and a slightly higher fraction of positive ratings (99% vs 98%).

Nevertheless, the Buybox is actually dominated by B: B has a Buybox share of 69.10%. This advantage emerges due to B's faster reaction time. In particular, the average reaction time for B is 34.74 minutes, compared to A's 56.63 minutes. The average cycle takes 24.11 hours and has an amplitude of $0.19. While this might seem small, the cycling itself indicates the potential presence of supra-competitive prices: the welfare loss we are concerned with is not that of price volatility but that of high prices in general. In particular, our model below will suggest that cycles are not anchored at marginal cost.

In this vein, we can focus on B (a seller employing the repricing company's services) and discuss its potential profit margin. To this end, note from the top panel of Figure 3 that prices dramatically dropped when C entered on 2018-10-09. Assuming that B's costs do not suddenly change when C enters, we can conservatively estimate its total costs (including fees) as bounded above by the price that it chooses (on average) after entry: $23.26. If these were B's marginal costs, the implied margin at the average prices during the depicted cycling period would be 31.39%.[7] This margin might not seem extraordinary, but note that the seller competes directly with another seller for this exact product. Thus, we would expect high residual price elasticities (as both sellers sell the t-shirt in the same brand, size, and color) and hence low profit in a static Bertrand-Nash equilibrium.

3 Descriptive Evidence

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 price. For each of the twenty lowest-priced offers[8] on a product, they contain information on the offer's listing price, shipping cost, and whether it is in the Buybox; we sum listing price and shipping cost and use the resulting landed price in all analyses below. Furthermore, the notifications provide details on shipping: availability to ship immediately, the shipping time (minimum and maximum), and whether the offer is fulfilled by Amazon. Finally, they contain information about the seller: the seller's id, whether they are a 'Featured Merchant', their positive feedback percentage, and their total feedback count.

Our primary data source is the near-complete set of notifications received by the repricing company between 08/26/2018 and 03/25/2020. We subset to those notifications covering products for which we can unambiguously infer the repricing status of the offer registered with the company, which eliminates notifications between 08/26/2018 and 04/01/2019 when repricing status was not recorded. Furthermore, we drop all products on which Amazon has an offer. This leaves us with notifications covering 04/01/2019 to 03/08/2020 with 341 unique dates, 102,849 unique products, 347 unique source merchants[9], and 55,121 unique merchants; we provide information about this dataset (aggregated to the daily level) in Table 1. For each offer (source and non-source), we observe prices at infinite[10] time resolution (which we aggregate to the daily level, e.g., the daily mean or minimum price). Furthermore, whenever there is a price change, we observe which offer currently holds the Buybox (which we aggregate to measure a seller's Buybox daily share which is the fraction of the day that the seller is in the Buybox.). For the repricing companies' clients only, we also observe whether the repricer is enabled and whether a resetting strategy is being used. Finally, for a subset

[7] Arguably, this margin is an underestimate: it assumes B is selling at cost post-entry, but B only reduces the prices just enough to beat C. Furthermore, B's self-reported costs plus fees lie at $13.10. Indeed, the same t-shirt sold (on Amazon) for as low as $18.49 in 2017 and $9.99 in 2019.

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

[9] We observe a product if and only if one of the repricing companies' client merchants sells it; we refer to these merchants as the 'source' merchants.

[10] That is, we observe any price change, no matter how brief.

	N	Mean	Std. Dev.	Min.	5%	50%	95%	Max.
Offer-Dates								
Mean Buybox Share	37,256,297	0.20	0.38	0.00	0.00	0.00	1.00	1.00
# Unique Prices	37,256,297	2.05	8.18	1.00	1.00	1.00	3.00	3,099.00
Mean Price	37,256,297	42.92	50.19	0.01	6.97	25.95	137.43	999.99
Max. Price	37,256,297	43.79	211.97	0.01	6.99	26.05	138.92	1,000,022.18
Min. Price	37,256,297	42.75	49.95	0.00	6.95	25.84	136.90	999.99
Product-Dates								
Mean Buybox Price	6,766,111	40.40	51.51	0.01	5.95	22.95	138.00	999.00
Mean Min. Price	7,922,130	43.48	55.25	0.16	5.95	24.74	148.09	999.00
Quantity Shipped*	2,187,594	0.34	1.87	0.00	0.00	0.00	2.00	312.00
Reported Cost (incl. Fees)*	1,933,246	25.71	24.12	1.49	6.64	18.26	73.25	860.40
Repricer Enabled?*	7,922,130	0.31	0.46	0.00	0.00	0.00	1.00	1.00
Resetting Active?*	1,492,682	0.04	0.18	0.00	0.00	0.00	0.00	1.00

Table 1: Summary Statistics.

Notes: This table provides summary statistics for the dataset underlying our analysis. For the first panel, the unit of observation is an offer on a given date; the second panel aggregates to the product by date level (one product consists of multiple offers by various merchants). *Starred variables are only available for the merchant registered with the repricing company. As there is one such source merchant per product, we report these variables in the product-date panel. *Quantity Shipped* and *Reported Cost* are only available for a subset of repricing merchants. The dataset covers data from 2019-04-01 to 2020-03-08 with 341 unique dates, 102,849 unique products, 347 unique source merchants and 55,121 unique merchants. The data has been filtered to exclude all products on which Amazon has an offer, and to exclude product-dates with mean prices above 999.99 USD (likely errors).

of these clients, we observe the quantity shipped and the reported cost of the item.

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

Our model below will assume a high amount of turnover in the market. This assumption is also supported by the data: following the average product over time, Figure 4a shows that the fraction of initial merchants still present declines quickly, and yet the total number of merchants per product increases quickly, indicating that both exit and entry are common phenomena.

There are also drawbacks to the data we employ. In particular, both the 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, as illustrated in Figure 4b, the number of *competitive* offers is typically far lower; and as 87.92% of uncompetitively priced offers are passively priced (using the definitions from Figure 4b), they often simply present a ceiling above which no algorithm can price.

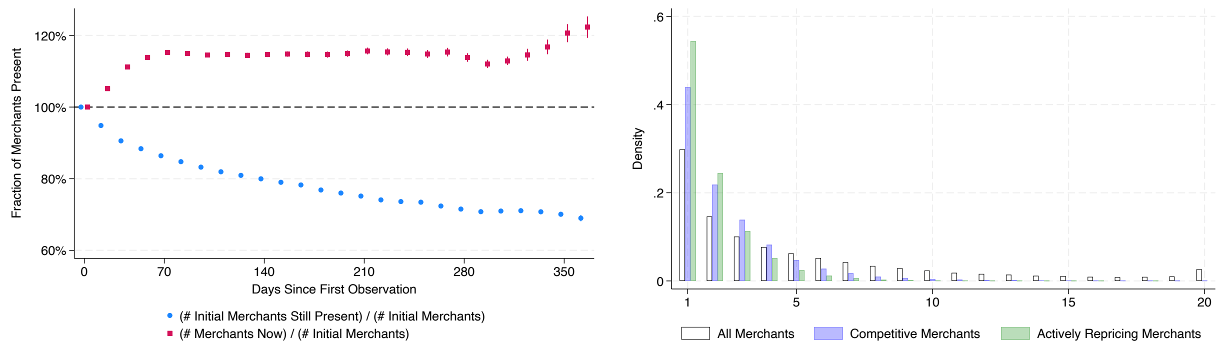

(a) Evidence on Turnover.

(b) Seller Histogram.

Figure 4: Evidence of Turnover and Seller Histogram.

Notes: The left panel investigates, after a given number of days following the first observation of a product, (i) what fraction of initial merchants (blue circles) are still present and (ii) the ratio of current to initial merchant count (red squares.) We aggregate to product-dates and regress each outcome on offer level fixed-effects and relative time fixed-effects. Standard errors are clustered at the product level. Without turnover, the blue circles and red squares would stick close to the horizontal line at 100%. Instead, they diverge. Indeed, only 56% of a merchant's competitors are those he faced a year ago. The right panel exhibits the distribution of the number of merchants across products. We consider an offer 'competitive' if its price is no more than 5% above the Buybox price and 'actively repricing' if at least two distinct prices exist for that offer on a given date. The modal number of offers is two for products with any competition, and the distribution becomes more concentrated near this mode when only considering competitive or actively repricing offers.

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.io) or "powered by AI" (goaura.com). However, a closer examination of repricing interfaces suggests that this wide variety in marketing claims is not reflected in the sophistication of implementable strategies.

We exhibit in Figure 5 a screenshot of a typical repricing interface; Appendix F contains additional examples to show that the interfaces do not vary much in implementable strategies (even when they vary in design.) Essentially all interfaces proceed by allowing merchants to set the amount by which their own offer's price should undercut (or overcut) either the currently lowest-priced competitor or the Buybox winner, with potentially varying behaviour against (e.g.) competitors fulfilled by Amazon (FBA). Furthermore, many (but not all) interfaces allow merchants to specify the repricing behaviour if the optimally calculated price were to fall below a pre-specified minimum[11]. This allows implementation of a cycling strategy by resetting prices when a competitor prices below your minimum. Finally, many repricing interfaces allow merchants to turn off repricing and reset prices to some maximum at a pre-specified time, once again allowing for implementation of a cycling strategy.

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 (e.g. Figure 5 but also the other interfaces in Appendix F) are not consistent with this interpretation.

[11]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.

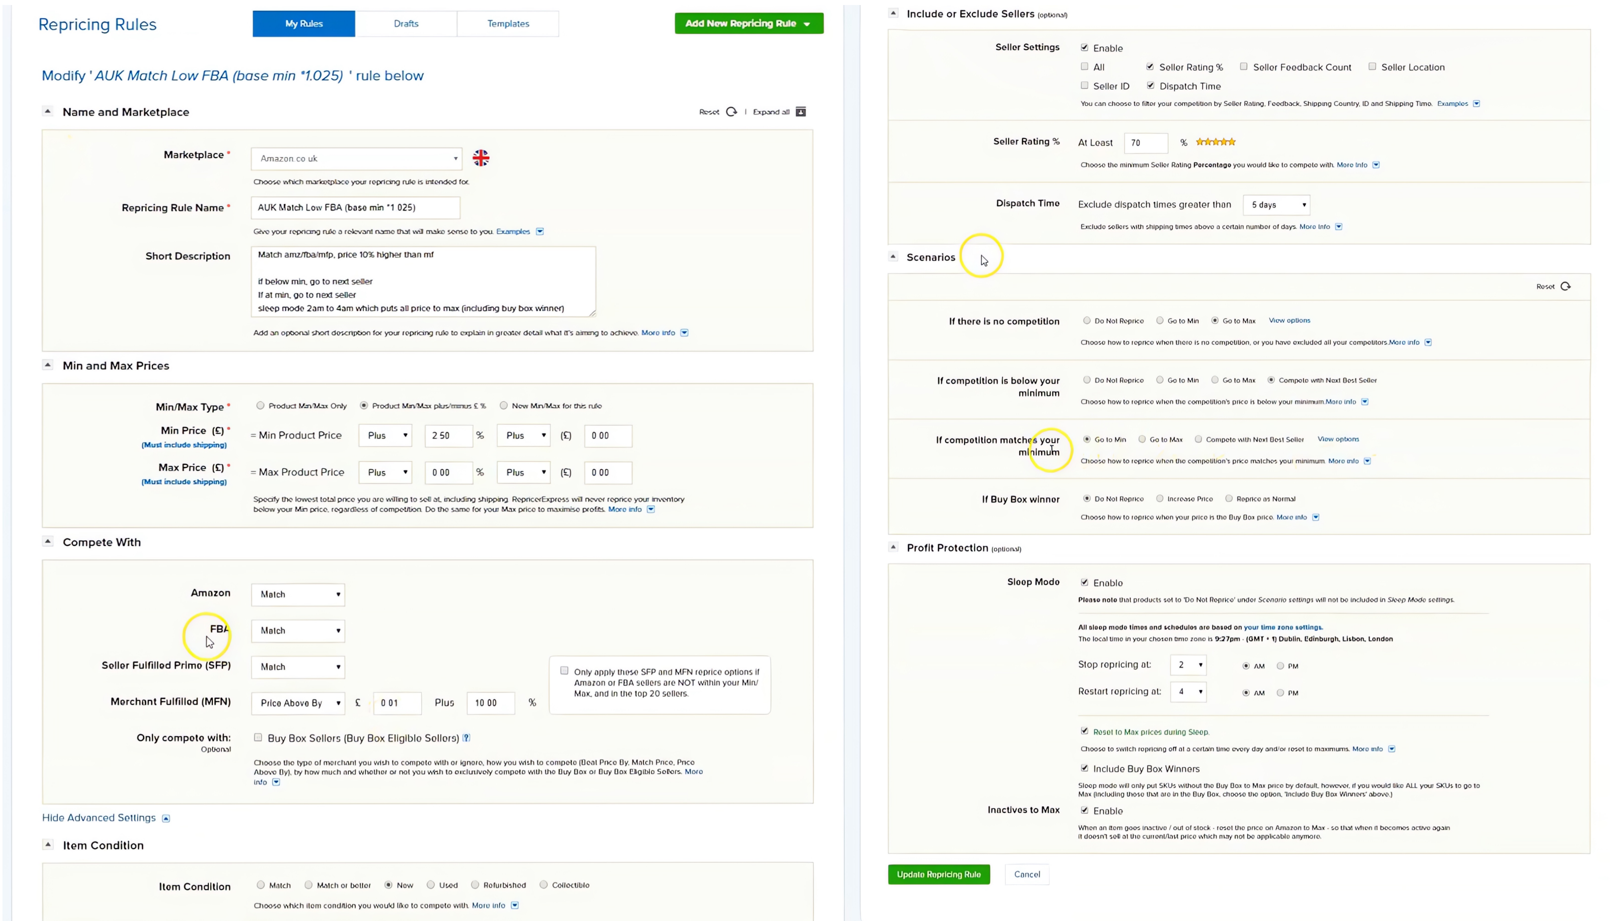

Figure 5: Repricing Interface Example (RepricerExpress.com).

Notes: This figure depicts the repricing interface for custom strategy creation at RepricerExpress. This interface is representative of those used by most repricers: it allows merchants to set the amount by which their offer's price should undercut (or overcut) each competitor, with potentially varying behaviour against competitors based on whether they are fulfilled by Amazon (FBA). Furthermore, it allows merchants to specify the repricing behaviour if the optimally calculated price falls below a pre-specified minimum (which varies offer by offer.) Finally, it has a 'Sleep Mode' (in the 'Profit Protection' section) that allows merchants to turn off repricing and reset prices to some maximum at a pre-specified time every day.

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 must randomize between resetting prices or keeping them unchanged, hoping that their rival might be the one to reset prices. However, neither randomization nor jumps are implementable using the repricing software interfaces.

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

Repricer Activation Event Study. To identify a causal effect of repricer activation on the outcomes of the offer for which repricing was activated, we exploit that signing up for the repricer and turning it on are separate actions. Some merchants activate the repricer 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 an event study design around the activation of the repricer. As we only observe merchants after they sign up with the repricer, our treatment effects will be estimated relative to a control group of merchants who signed up but did not yet activate the repricer.

We start by building intuition for what happens when the repricer is activated. Figure 6 exhibits two examples of typical price paths after repricer activation. In both panels, the horizontal axis measures days since the repricer was activated and the vertical axis measures price. The orange line corresponds to the price of the offer for which repricing was activated, and the blue line corresponds to the price of the lowest-priced offer amongst all other offers; the gray area indicates the period during which we cannot ascertain treatment status. In both panels, the repricer aggressively cuts prices when it is activated. However, while this price cut leads to no response in the left panel, it triggers a price war that lasts for days in the right panel.

Are the price decreases in these examples representative of how repricer activation affects prices? To answer this question, we estimate an event-study specification[12] of the form

$$y_{it} = \sum_{k=-L_G}^{L_M-1} \delta_k \Delta z_{i,t-k} + \delta_{L_M} z_{i,t-L_M} + \delta_{-L_G-1}(1 - z_{i,t+L_G}) + \alpha_i + \gamma_t + \epsilon_{it}, \tag{1}$$

where z_{it} is a dummy measuring whether offer i has repricing turned on at date t, Δ denotes the first difference operator, α_i are offer fixed-effects, γ_t are day fixed-effects, and ϵ_{it} is an error term. The coefficients of interest are the δ_k coefficients, which measure the cumulative effect of repricing k days after the initial activation (with k potentially negative). The specification implicitly assumes that the effect of activating the repricer stabilizes at δ_{L_M} after L_M days; similarly, it assumes that the pre-activation difference between control and treatment offers is constant at δ_{-L_G-1} before L_G days

[12]The specification we use is explained in more detail in Freyaldenhoven et al. (2021).

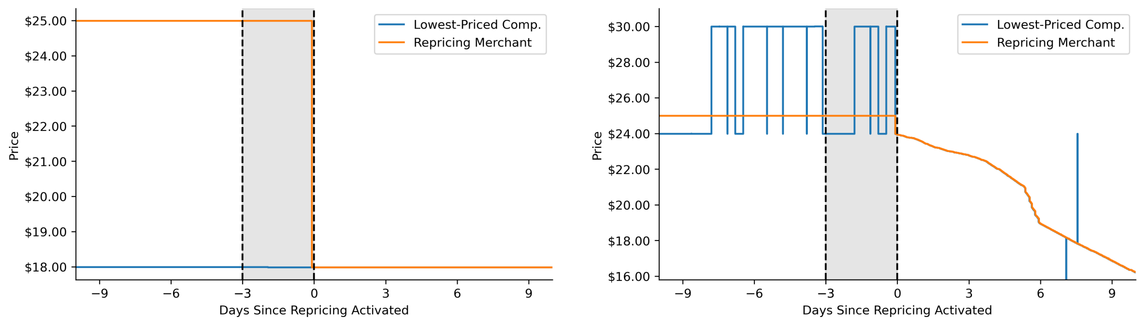

Figure 6: Examples of Price Changes After Repricer Activation.

Notes: This figure exhibits two examples of typical price paths after repricer activation. In both panels, the horizontal axis measures days since the repricer was activated and the vertical axis measures price. The orange line corresponds to the price of the offer for which repricing was activated, and the blue line corresponds to the price of the lowest-priced offer amongst all other offers (the identity of which can vary over time). The gray area indicates the period during which we cannot ascertain treatment status (due to sampling frequency limitations). In both panels, the repricer aggressively cuts prices when it is activated. However, while this price cut leads to no response in the left panel, it triggers a price war that lasts for days in the right panel.

preceding the activation event. Compared to a difference-in-differences specification, this event-study specification allows us to investigate a potentially dynamic treatment effect. Furthermore, it enables us to examine pre-trends to assess the plausibility of the identifying assumption of counterfactual parallel trends between treated and control units.

We exhibit our results in Figure 7 and (corresponding) Table 2. We begin by noting that there are essentially parallel pre-trends, and these are quite precisely estimated depending on the outcome (pricing outcomes are estimated with more precision). While the confidence intervals do not overlap with zero in all cases, they are small compared to the treatment effects. Thus, we conclude that the absence of selection into treatment (i.e., counterfactual parallel trends) is plausible.

We now move on to discussing the results. As expected, repricer activation leads to an instant, large, and statistically significant increase in the number of unique prices an offer takes on for any given day. However, interestingly, this effect dies down over time. While this pattern would naturally emerge if merchants that initially experimented with repricing abandoned it at increasing rates over time, we can rule this out given the stability of repricing usage after the initial adoption reported in Figure E.1. Instead, the evidence suggests that the initial flurry of repricing activity slows down because the repricer frequently finds that its optimal action, given the price adjustments by its rivals, is to leave the price unchanged, possibly because it undercuts competitor prices but eventually reaches the minimum price entered as a safeguard by the merchant.

Moving on to the third (middle-left) panel of Figure 7, we see that prices drop dramatically on activation of the repricer: the day after we are confident that the repricer has been turned on (at most three days after the repricer was turned on), prices are already down by 8%. This drop is persistent and deepens over time before stabilizing at about a 16.7% drop after 50 or so days. This pattern is consistent with the repricer trying to aggressively undercut opponents' prices to gain the Buybox but being faced with competitors doing the same. However, it is noteworthy how long the overall effect takes to play out: as repricers can change prices as often as every two minutes, the fact that it takes the repricer 50 days to reach its final price suggests that some competitors may be competing with the repricer manually.

While the repricing merchants' prices drop dramatically on the day repricing is activated, the fourth and fifth (middle-right and bottom-left) panels of Figure 7 show that the repricer is not dropping prices excessively upon impact. In particular, the mean minimum price across offers[13] and the Buybox price potentially increase on impact. While this may seem counterintuitive, it is consistent with undercutting or matching the lowest-priced competitor. If this competitor was previously above the now-repricing merchant's price, the repricer would increase the price to match the competitor's price. If he was below, the repricer would decrease the price – but just by enough to undercut the competitor. This minor price cut is not enough (on average) to offset the price increase when the competitor exceeds the repricer's price. Still, note that the positive effects on the minimum and Buybox prices do not last very long: indeed, we see the same indication of a prolonged price war here as in the repricing merchant's price. Nevertheless, as both the minimum and the Buybox price were, on average, lower than the repricing merchant's price before repricer activation, the overall decline e.g. in Buybox price stabilizes at only 9.5%.

The final question that remains to be answered[14] is what these price declines buy the repricing merchant. The sixth (bottom-right) panel indicates that the repricing merchant can dramatically increase its Buybox share after it begins repricing: on a basis of a 26% share of the Buybox, it gains about 9pp to reach a 35% share of the Buybox. These effects are large and persistent; about half of the effect is already present the day after activation, and the effect peaks after 11 days. While the effect somewhat decreases over time, it never falls below the initial effect.

Finally, there is a recent concern in the applied econometrics literature that two-way fixed effects estimators (such as the one employed for the results in our paper) perform poorly in the presence of treatment effect heterogeneity. Thus, we also estimate the same specification using the Sun and Abraham (2021) approach that is robust to these concerns; we exhibit the results in Figure E.2. The results are qualitatively similar to the ones we obtained using the two-way fixed effects estimator, and indeed, almost all event study coefficients are statistically indistinguishable from the ones we obtained using the two-way fixed effects estimator.

The Importance of Resetting Prices. When we see unusual patterns, it is often possible to explain the same patterns with various models: when it comes to price cycles, for instance, repricing strategies and consumer-search-based theories are both ex-ante possible explanations. Another possibility is that merchants prefer to make a small volume of sales at a positive margin than to make many sales at near-zero margins; thus, they may wait for prices to drift down to their marginal cost and then reset prices without anticipating that their competitors will follow suit. Crucially,

[13]This is the time-weighted daily average of the minimum price, where the minimum across offers is taken first and the time-weighted average over time second.

[14]Ideally, we would be interested in effects not just on Buybox share, but effects on sales and profits. These are rare events, and hence estimated with much noise; furthermore, only a small sample of offers reports cost information, and this sample could be selected. We hence discuss these effects in Appendix E.

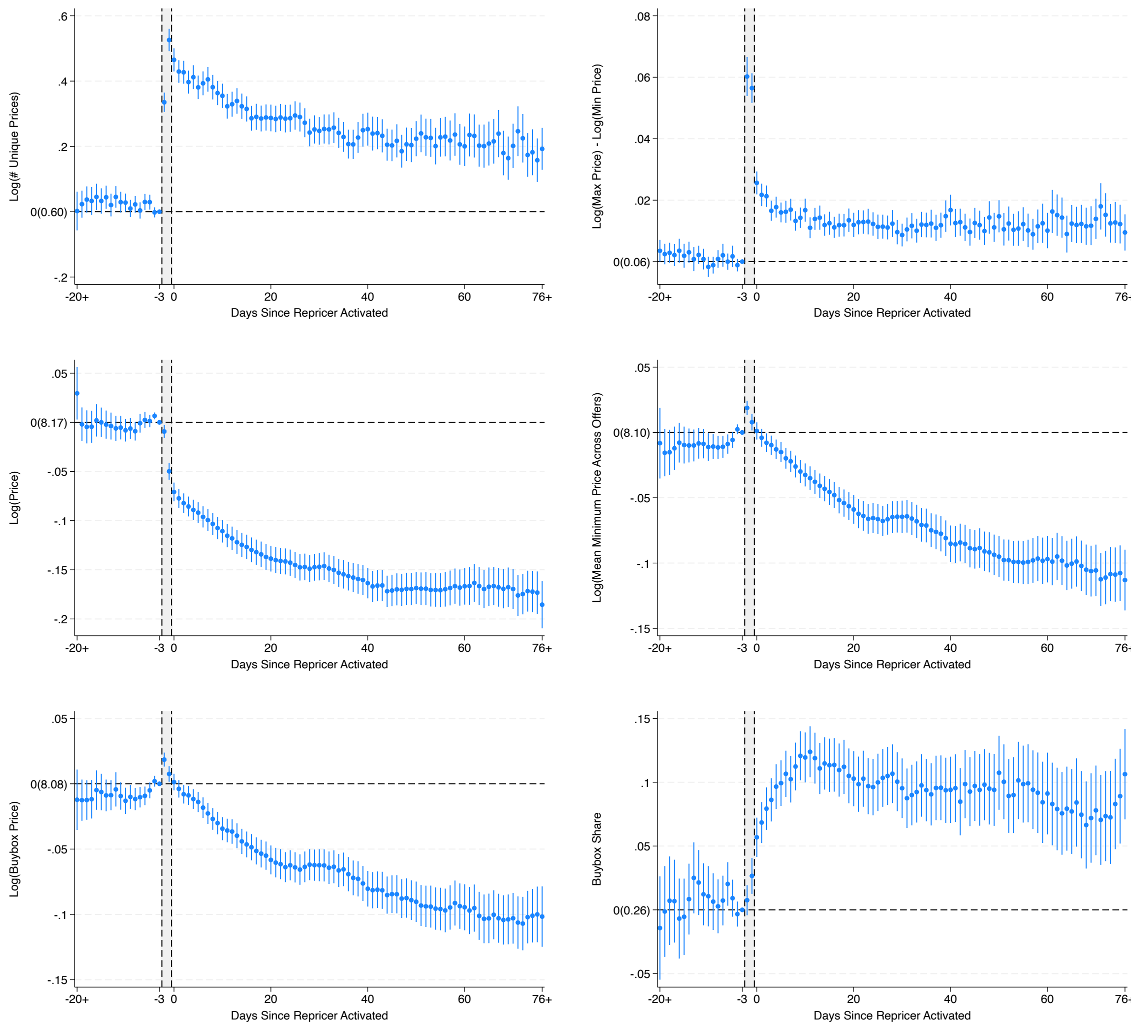

Figure 7: Activating the Repricer Lowers Prices and Increases Buybox Share.

Notes: This figure exhibits results from an event study that measures the effects of initial repricer activation. The horizontal axis measures days since the repricer was activated (with 0 being the day of activation.) Due to sampling frequency limitations, repricing status is unknown during the time window indicated by light gray shading; thus, we normalize the coefficient that measures the treatment effect three days before (certain) treatment to zero to avoid expressing treatment effects relative to a period where treatment has already started for some units. The vertical axis measures the effect of activating the repricer on the outcome variable of interest, with a zero value indicating no effect; we also provide the mean of the outcome variable three days before treatment in the parenthetical label. Each blue dot corresponds to a coefficient δ_s in (1), and the bars indicate a 95% confidence interval derived from standard errors clustered at the offer-level. '# Unique Prices' refer to the total number of prices any offer takes on on a given day. 'Max Price' and 'Min Price' refer to the daily maximum and minimum price for a given offer (i.e., the minimum and maximum are taken across time within offer.) 'Price' refers to the time-weighted average price of an offer on a given day. The 'Mean Minimum Price Across Offers' refers to the time-weighted daily average of the minimum price, where the minimum across offers is taken first and the average over time second. 'Buybox Price' refers to the time-weighted daily average Buybox price. 'Buybox Share' refers to the fraction of time the offer is in the Buybox.

	(1) Log # Unique Prices	(2) Log Price Span	(3) Log Price	(4) Log Mean Min.	(5) Log Buybox Price	(6) Buybox Share
Days Since Repricer Activ.						
-15 Days	0.033*	0.002	-0.000	-0.010	-0.006	-0.005
	(0.020)	(0.002)	(0.008)	(0.007)	(0.007)	(0.015)
-5 Days	0.029**	0.002	0.001	-0.006*	-0.005	0.009
	(0.012)	(0.002)	(0.003)	(0.003)	(0.003)	(0.007)
0 Days	0.465***	0.026***	-0.071***	0.001	0.001	0.057***
	(0.018)	(0.002)	(0.005)	(0.003)	(0.003)	(0.008)
5 Days	0.381***	0.016***	-0.092***	-0.015***	-0.014***	0.099***
	(0.018)	(0.002)	(0.005)	(0.004)	(0.004)	(0.009)
15 Days	0.315***	0.012***	-0.127***	-0.046***	-0.046***	0.113***
	(0.020)	(0.002)	(0.006)	(0.005)	(0.005)	(0.011)
30 Days	0.247***	0.009***	-0.147***	-0.065***	-0.062***	0.095***
	(0.022)	(0.002)	(0.007)	(0.006)	(0.006)	(0.012)
60 Days	0.200***	0.010***	-0.167***	-0.097***	-0.095***	0.091***
	(0.031)	(0.003)	(0.010)	(0.009)	(0.009)	(0.016)
Omitted Coef.	✓	✓	✓	✓	✓	✓
Offer FE	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
# Clusters	27,314	27,314	27,314	27,314	25,532	27,314
# Observations	742,968	742,968	742,968	742,968	662,769	742,968
Mean Dep.	0.60	0.06	8.17	8.10	8.08	0.26

Table 2: Event-Study Estimates of Repricer Activation Effect.

Notes: This table provides event-study estimates of the effects of repricer activation based on equation (1) and corresponding to (a subset of the) coefficients plotted in Figure 7. Due to the large number of event-study coefficients, we exhibit a representative subset; the underlying regression includes all omitted event-study coefficients from -20+ to 76+. Each regression is run on a a potentially unbalanced panel dataset at the *offer x date* level and includes offer and date fixed-effects. The panel is not necessarily balanced because offers can disappear (e.g., if they sell out) and reappear (e.g., if they acquire additional inventory) over time. The sample size changes for (5) as a Buybox Price does not exist when no offer is in the Buybox. Standard errors clustered at the offer level in parentheses; *** p<0.01, ** p<0.05, * p<0.1. '# Unique Prices' refer to the total number of prices any offer takes on on a given day. The 'Price Span' refers to the difference between the daily maximum and minimum price for a given offer (i.e., the minimum and maximum are taken across time within offer.) 'Price' refers to the time-weighted average price of an offer on a given day. The 'Mean Min.' refers to the time-weighted daily average of the minimum price, where the minimum across offers is taken first and the average over time second. 'Buybox Price' refers to the time-weighted daily average Buybox price. 'Buybox Share' refers to the fraction of time the offer is in the Buybox.

however, this strategy makes sense only in settings where demand is not very elastic: in our setting, raising price by just 5% leads to a more than 75% drop in Buybox share.

To support the idea that price resets are *strategic*, i.e., that merchants expect their competitors to follow them up, we now provide some anecdotal evidence. In particular, repricing companies need to advertise their services, and they often do so by explaining to would-be clients why using their software could increase their profits. We now provide some examples of such advertising, some very explicit in its discussion of the strategic benefits of resetting.

To begin with, SellerEngine reminds sellers in a blog post that "as you reprice upwards, Amazon sellers who use automatic repricing software may follow your lead and raise prices as well" and hence "you can fight the trend of lower prices, by periodically raising your prices" (SellerEngine.com). At the same time, Aura's blog advocates for a strategy that "allows you to lower your price by $0.01 to increase your Buy Box percentage, which increases the volume of sales. The software then raises your price to your Maximum price once you've reached your Minimum, essentially increasing the average net profit" (goaura.com). Similarly, Informed features a 'Smart Price Reset' strategy (see Figure F.1) which it describes as "a smart way to essentially 'reset' prices on a listing to prevent [...] the dreaded race to the bottom" (informed.co), helpfully noting that this will cause "your competitors to raise their prices in response, thus helping to increase your profit" (informed.co). Not to be left behind, RepricerExpress also offers a 'Sleep Mode' strategy (see Figure 5), pointing out 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). This is consistent with Repricer.com's reasoning, which suggests that its 'Pause repricing' feature can "halt repricing during periods of low sales with the expectation that other sellers employing a Repricer will follow suit, thereby avoiding a price war" (repricer.com). Just in case there was any doubt as to why all of these 'resetting' strategies are advantageous, a (since deleted) comment on popular forum Quora.com puts a rather fine point on it by declaring that "Colluding with other repricers for night bumps is very profitable" (Figure F.9).

While these strategies are straightforward, more complex cycling strategies are also present and go by such names as "yo-yo" pricing. For instance, SellerSnap's homepage features a description of its 'AI Algorithmic Repricer', which they suggest is equipped to deal with aggressive competitors because it "will mirror the behavior of the competitor by reducing the price until the reset point. [...] When this happens, the repricer will increase the price to your maximum" (sellersnap.io). SellerSnap even confirms that they understand the tradeoffs of resetting by noting that "You will temporarily lose the Buy Box while waiting for your competitor to raise the price" (sellersnap.io). Still, if you are not ready to jump on the AI bandwagon quite yet, SellerSnap offers a 'Yo-Yo Repricing Rule' (see Figure F.8) that "allows you to manually set a loop increaseing the price to Max and then reverting to an automatic repricing method" (sellersnap.io). Of course, this rule is not only a feature at SellerSnap – ChannelMAX, for instance, also offers what it calls the 'Amazon YoYo', advocating for its use explicitly when competing with Amazon itself as "It's nearly impossible to get BuyBox when Amazon is selling. But if you can make Amazon go higher and come down at a faster pace,

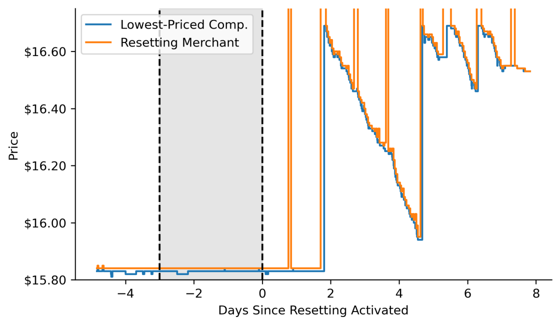

Figure 8: Examples of Price Changes After Resetting.
Notes: This figure exhibits an example price path after a merchant activates a resetting strategy. The horizontal axis measures days since the repricer was activated and the vertical axis measures price. The orange line corresponds to the price of the offer for which resetting was activated, and the blue line corresponds to the price of the lowest-priced offer amongst all other offers (the identity of which can vary over time). The gray area indicates the period during which we cannot ascertain treatment status. When resetting is activated, the competitor responds by following the resetting merchant up.

then you surely can get some BuyBox" (channelmax.net) and "if this process is repeated over and over again, you surely can benefit" (channelmax.net).

The above discussion suggests that merchants may want to increase prices periodically to avoid price wars from eroding profits. We refer to this behavior as resetting and now ask: Does resetting work, i.e., do opponent prices increase in response to a merchant resetting? To answer this question, we exploit the fact that we observe the prices of all offers in the marketplace, not just those of the merchant that resets. We can thus estimate the effect of resetting on competitors' prices.

Before moving on to our key results, we warm up by considering the example price path in Figure 8. The figure illustrates the experience of a single client of the repricer who activated a resetting strategy. His price (orange) is stable before the strategy is activated and regularly spikes after the strategy is activated. However, note how his lowest-priced competitor (blue) already makes (minimal) price changes in the pre-period. These changes indicate that the competitor is actively repricing and, hence, may be able to follow a price reset quickly. Indeed, when the price reset is activated, the competitor responds by following the resetting merchant up.

Moving on from our example to broader results, we estimate regressions of the form

$$y_{it} = \sum_{k=-L_G}^{L_M-1} \delta_k \Delta z_{i,t-k} + \delta_{L_M} z_{i,t-L_M} + \delta_{-L_G-1}(1 - z_{i,t+L_G}) + \alpha_i + \gamma_t + \epsilon_{it}, \tag{2}$$

where z_{it} is a dummy measuring whether offer i has resetting turned on at time period t, Δ denotes the first difference operator, α_i are offer fixed-effects, γ_t are time fixed-effects, and ϵ_{it} is an error term. In contrast to the activation event study, we aggregate to ten-day periods to tamp down noise. The coefficients of interest are the δ_k coefficients, which measure the cumulative effect of resetting strategies k periods after the initial activation (with k potentially negative). The specification implicitly assumes that the effect of resetting stabilizes at δ_{L_M} after L_M periods; similarly, it assumes that the pre-activation difference between control and treatment offers is constant at δ_{-L_G-1} before L_G periods preceding the resetting activation event. Compared to a difference-in-differences specification, this event-study specification allows us to investigate a potentially dynamic treatment effect. Furthermore, it enables us to examine pre-trends to assess the plausibility of the identifying assumption of counterfactual parallel trends between treated and control units. In contrast to the repricing activation event study, when activating resetting, most merchants have already been with the repricer for a while; hence, we estimate more pre-treatment coefficients.

To give ourselves the best chances of isolating the actual effect of resetting, we only keep data from products that (on average) have between two and five competitive offers: with a single competitive offer, we would expect no effect of resetting; and if there are many competitive offers, coordination on resetting is also challenging.

We exhibit the results from this specification in Figure 9 and Table 3. As in our other event study, we see parallel pre-trends with the possible exception of the number of unique prices, which exhibits a slight upward trajectory that is small compared to the estimated treatment effect. Hence, we conclude that the absence of selection into treatment on time trends (i.e., counterfactual parallel trends) is plausible. Furthermore, the Sun and Abraham (2021) estimates essentially coincide with the two-way fixed effects estimates.

The first row of Figure 9 acts as a first stage: Does the pricing behavior of the focal merchant change when he activates the resetting strategy according to the repricing company? The answer is a resounding yes: the repricing offer's daily number of unique prices essentially doubles, and the repricing offer's daily maximum price increases by around 10% on impact.

Moving on to the third panel (the first column in the second row), we can see that the focal merchant's increase in maximum price is eventually met by a similar increase in maximum price by his competitors. However, this effect is somewhat delayed and takes more than a month to kick in, suggesting that it perhaps requires human intervention. By contrast, the focal merchant's daily minimum price (fourth panel) starts to increase smoothly immediately after he adopts the resetting strategy. One possible explanation for this divergence is that the competitor effects are averaged across multiple competitors, and only the most aggressive competitor may be of immediate relevance to the focal merchant's pricing decision. Indeed, the final two panels investigate the effect on market prices (i.e., the Buybox price and the minimum price across offers) and find that these increase relatively quickly after the focal merchant adopts the resetting strategy.

Finally, Figure E.6 in the appendix shows that activation of the resetting strategy is associated with increased profits for the (possibly selected) sample of offers for which merchants report cost information. We caution that due to the rare nature of sales on the long tail of Amazon, these estimates are very noisy and do not reach conventional levels of statistical significance. Furthermore, we have argued above that the sample for which we have cost information is likely selected, and hence these results should be taken with a grain of salt.

Having established that resetting can successfully raise prices, we want to know whether merchants think strategically about when to increase prices. In Appendix D, we develop an algorithm to

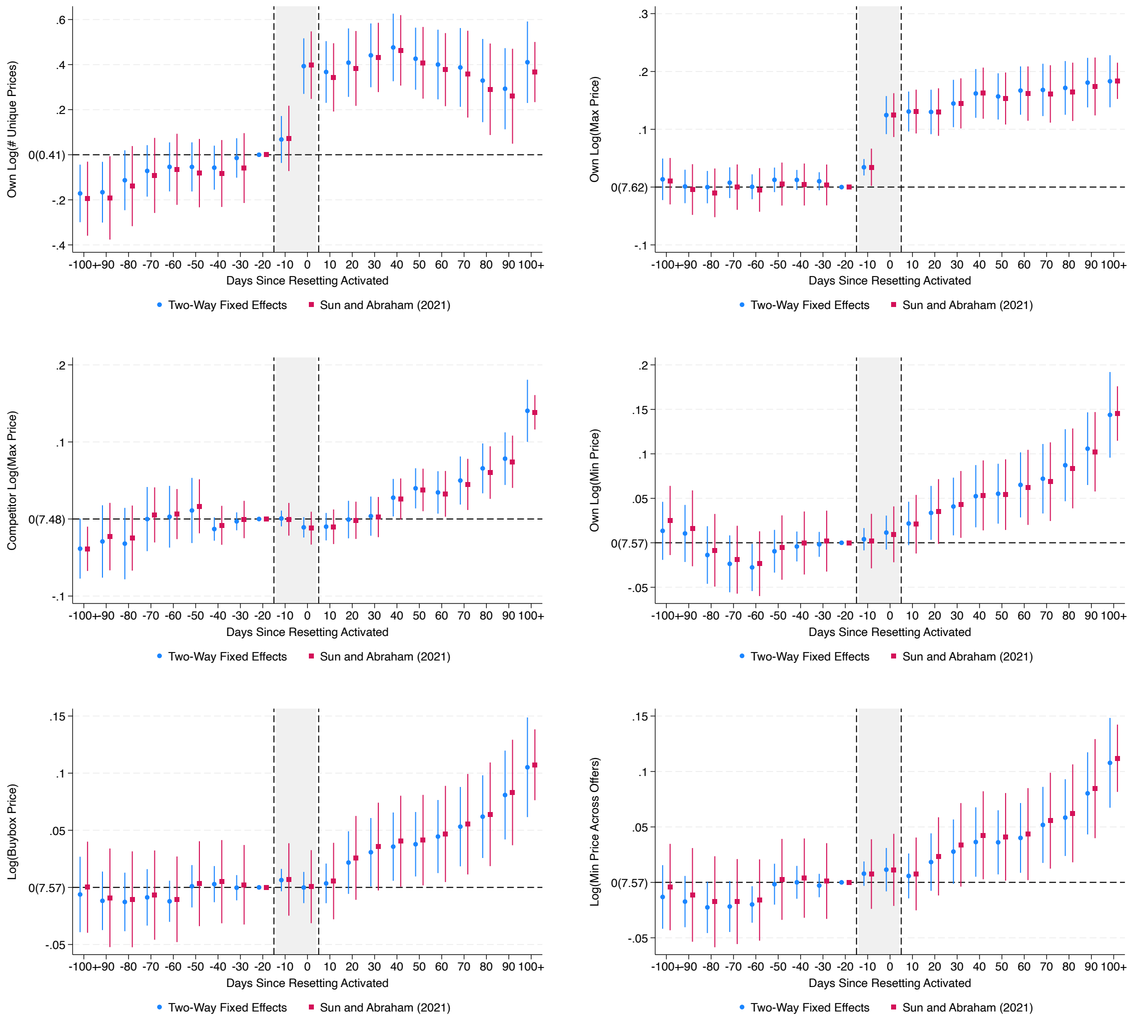

Figure 9: Activating Resetting Raises Prices.

Notes: This figure exhibits results from an event study that measures the effects of activating a resetting strategy. The horizontal axis measures days since the resetting strategy was activated for the source merchant's offer (with 0 being the day of activation.) Due to sampling frequency limitations, resetting status is unknown during the time window indicated by light gray shading; thus, we normalize the coefficient that measures the treatment effect twenty days before (certain) treatment to zero to avoid expressing treatment effects relative to a period where treatment has already started for some units. The vertical axis measures the effect of activating the resetting strategy on the outcome variable of interest, with a zero value indicating no effect; we also provide the mean of the outcome variable twenty days before treatment in the parenthetical label. Note that outcome variables could be associated with the offer that started employing the resetting strategy (e.g., "Own Log(# Unique Prices)"), with the outcomes for competitors (e.g., "Competitor Log(Min Price)"), or with market-level outcomes (e.g., "Log Buybox Price".) Each blue dot corresponds to a coefficient β_s in (2), and the bars indicate a 95% confidence interval derived from standard errors clustered at the offer-level. '# Unique Prices' refer to the total number of prices any offer takes on on a given day. 'Max Price' and 'Min Price' refer to the daily maximum and minimum price for a given offer (i.e., the minimum and maximum are taken across time within offer.) 'Buybox Price' refers to the time-weighted daily average Buybox price. 'Min Price Across Offers' refers to the time-weighted daily average of the minimum price, where the minimum across offers is taken first and the average over time second.

	(1) Own Log # Unique Prices	(2) Own Log Max Price	(3) Comp. Log Max Price	(4) Own Log Min Price	(5) Market Log Buyb. Price	(6) Market Log Min Price
Days Since Resetting Activ.						
-90 Days	-0.166**	0.001	-0.029	0.010	-0.012	-0.017
	(0.069)	(0.015)	(0.024)	(0.016)	(0.013)	(0.012)
-60 Days	-0.054	0.001	0.003	-0.028**	-0.012	-0.020**
	(0.055)	(0.011)	(0.020)	(0.014)	(0.009)	(0.008)
-30 Days	-0.014	0.010	-0.003	-0.002	-0.000	-0.003
	(0.045)	(0.008)	(0.006)	(0.007)	(0.006)	(0.005)
0 Days	0.393***	0.125***	-0.011	0.011	-0.000	0.011
	(0.063)	(0.017)	(0.007)	(0.010)	(0.007)	(0.010)
30 Days	0.441***	0.145***	0.004	0.041**	0.031**	0.028*
	(0.072)	(0.021)	(0.013)	(0.017)	(0.015)	(0.015)
60 Days	0.400***	0.167***	0.034**	0.065***	0.045***	0.040**
	(0.079)	(0.021)	(0.014)	(0.019)	(0.016)	(0.016)
90 Days	0.293***	0.181***	0.078***	0.106***	0.081***	0.080***
	(0.092)	(0.022)	(0.017)	(0.021)	(0.020)	(0.019)
Omitted Coef.	✓	✓	✓	✓	✓	✓
Offer FE	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
# Clusters	8,114	8,114	11,051	8,114	7,853	8,114
# Observations	73,096	73,096	171,218	73,096	70,258	73,096
Mean Dep.	0.41	7.62	7.48	7.57	7.57	7.57

Table 3: Event-Study Estimates of Resetting Effect.

Notes: This table provides event-study estimates of the effects of resetting strategies based on equation (2) and corresponding to (a subset of the) coefficients plotted in Figure 9. Due to the large number of event-study coefficients, we just exhibit a representative subset of them here; however, the underlying regression includes all omitted event-study coefficients from -100+ to 100+. Each regression is run on a a potentially unbalanced panel dataset at the *offer x (10-days)* level and includes offer and time fixed-effects. The panel is not necessarily balanced because offers can disappear (e.g., if they sell out) and reappear (e.g., if they acquire additional inventory) over time. Standard errors clustered at the offer level in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Note that outcome variables could be associated with the offer that started employing the resetting strategy (e.g., "Own Log(# Unique Prices)"), with the outcomes for competitors (e.g., "Competitor Log(Min Price)"), or with market-level outcomes (e.g., "Log Buybox Price".) '# Unique Prices' refer to the total number of prices any offer takes on on a given day. 'Max Price' and 'Min Price' refer to the daily maximum and minimum price for a given offer (i.e., the minimum and maximum are taken across time within offer.) 'Buybox Price' refers to the time-weighted daily average Buybox price. 'Min Price Across Offers' refers to the time-weighted daily average of the minimum price, where the minimum across offers is taken first and the average over time second.

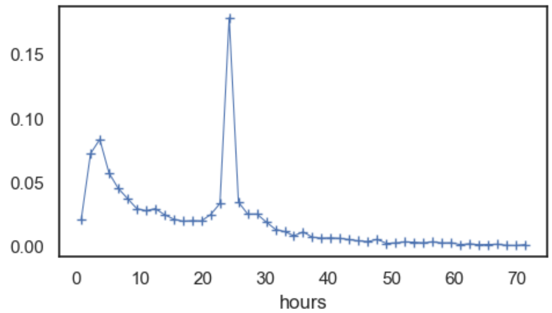

(a) Distribution of Cycling Period

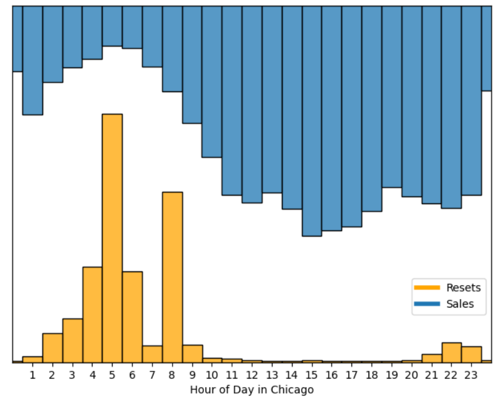

(b) The Reset Times of Daily Cycles

Figure 10: Day-Long Cycles Are Common, And Reset at Night.
Notes: The left panel shows the distribution of average cycling periods (i.e., the length of time between resets) for all products. This right panel shows the distribution of the reset times of 105,895 day-long cycles (in orange; measured from bottom) and the distribution of 8,866,256 sales (in blue; measured from top) by hour of day. We define a day-long cycle as a cycle that takes 24h $\pm 1h$ from peak to peak. The figure illustrates that day-long cycles are almost exclusively reset at night when the overall sales activity is lowest.

identify cycling offers. 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 demand over time. In particular, sales generally happen when people are awake. The fact that night-time sales are unlikely 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 the most frequent cycling period is daily as shown in Figure 10a. Furthermore, resets of these daily cycles occur during the night hours when sales probabilities are lowest. Figure 10b depicts histograms of sales (in blue, measured from the top) and reset times[15] (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.

Summary Statistics. While cycling is an intriguing phenomenon, it is important to emphasize that it is nevertheless *rare* in our data. Figure 11 shows the prevalence of cycling by product category. We see that, averaged across all offers for a product, and across all products, the average offer spends

[15] 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.

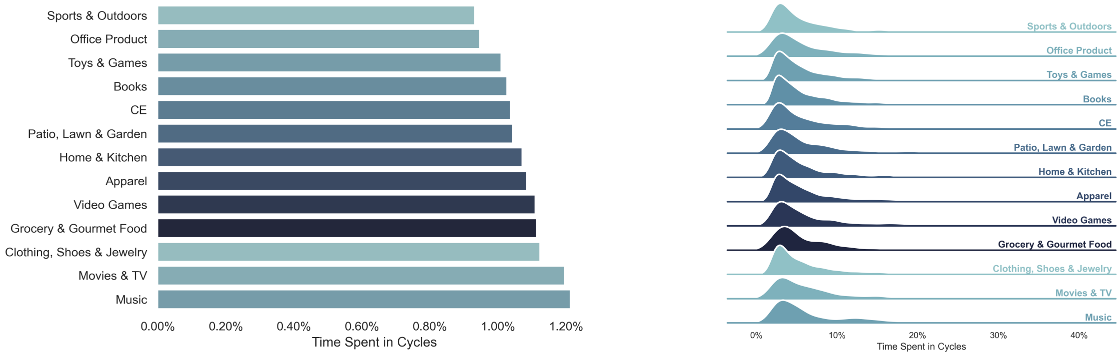

(a) Average Fraction Time Spent Cycling

(b) Distribution Conditional On Not Zero

Figure 11: 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.

1.07% of the time cycling. While this is low, several caveats are in order. Firstly, as discussed in the context of Figure 4b, many offers are uncompetitive, and even more are passive (i.e., rarely change prices) – cycling is more common amongst competitive and actively-repricing offers. Secondly, our cycle detection algorithm (discussed in more detail in Appendix D) is tuned to achieve a low false positive rate; thus, we will inevitably miss some cycling in our data.

For products with at least one cycle recognized, Figure 10a depicts the distribution of the average cycle length in hours. As discussed above, many products' cycles take 24h on average. Furthermore, Figure G.2 shows the distribution of the average amplitude as a fraction of an offer's mean price. The median (cycling) product with day-long cycles has an average amplitude of 8.02%, and the median other product has an amplitude of 4.74%. Consistent with the low resetting cost on Amazon, these amplitudes are somewhat smaller than those documented for other industries. For comparison, the typical amplitude in gasoline markets is 10%-13% (Noel, 2018, 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. These are the strategies that are implementable on the repricing interfaces, which are surprisingly uniform across different repricing companies (Appendix F exhibits examples from five different repricers.) Why this uniformity? Intuitively, the market started with all agents using a passive pricing technology. When the ability to reprice was introduced into the market, agents soon discovered that an undercutting strategy was a best-response to passive competitors. The cycling (or 'undercut then reset') strategy, in turn, is a best response to an undercutting strategy (Schlosser and Boissier, 2017).

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 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 $c_U = 0$. Still, such declining costs would only speed up the convergence to the equilibrium we describe below.

After 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 B and do not lead to supra-competitive profits. While this is immediately obvious for the (U, U) pairing, it is also true for (C, C). There is an incentive for the non-resetting merchant to set a fixed price below the cycle (which is achievable by setting $\underline{p}_C = \overline{p}_C$), thereby forcing his opponent to constantly price at the opponent's 'reset' level and capturing all demand himself. This incentive does not exist if the other merchant has chosen a very low $\underline{p}_C$, but in that case, the first merchant will prefer to reset the cycle himself to avoid prices drifting too far from monopoly price. In doing so, he in turn creates an incentive for the other merchant to set a fixed price below the cycle.

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. Two merchants produce output at zero marginal cost and face industry demand $D(p) = 20 - p$. As a first-order approximation of the 'Buybox' mechanism, we assume that only the currently lower-priced seller is exposed to this industry demand; the other seller faces zero demand[16].

[16] To be clear, this assumption is not descriptively accurate as (i) not all consumers purchase through the Buybox and (ii) the Buybox is not infinitely price elastic but has a price elasticity of around -20 (Lee and Musolff, 2023). 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. We investigate the robustness of our main results to this assumption in Appendix C, showing that they remain qualitatively similar for large but finite price elasticities.

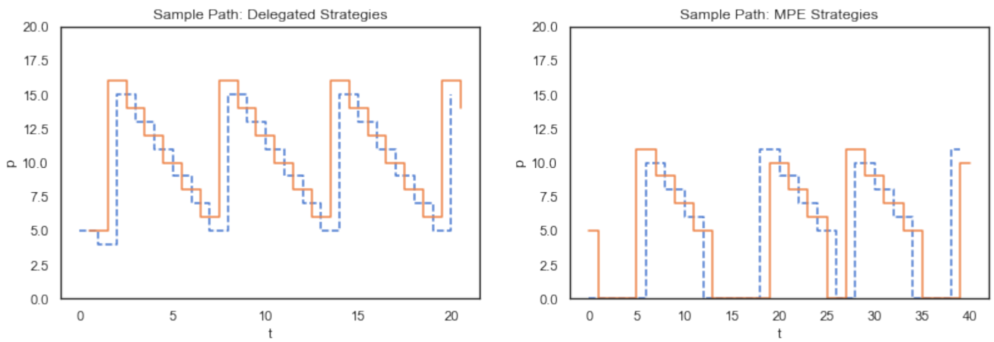

Figure 12: Sample Paths under MPE vs Delegated Strategies.
Notes: This figure provides example price paths for cycling under delegated strategies (left panel) and Markov-perfect strategies (i.e., the Maskin-Tirole equilibrium).

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.

Maskin-Tirole strategies are response functions $r_i(p_j)$, which specify agents' prices as a function of the last price set by their 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 equilibrium parameter choices of the merchants are[17]

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

Furthermore, following Maskin and Tirole (1988) we can find an Edgeworth MPE in this setup. Figure 12 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, 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 repricing interface restrictions, there is simply no way for merchants to jump down to costs. Thus, the agents cannot punish opponents for pricing below the cycle; indeed, as mentioned above, this is why the unique equilibrium amongst C agents is to price at cost. However, if the opponent always undercuts, 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

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

monopoly profit[18], 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 currency unit k ensures that jumps characterize real price paths. But as these jumps are small, smooth price paths will be a good approximation for sufficiently small k.

However, proper convergence to the continuous approximation cannot be achieved by letting $k \to 0$ as that 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 Δ that passes between two subsequent undercutting steps tend to zero. Thus, 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 unit of time until he is allowed to reprice again.

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

$$V_r(\ell, u) = \frac{1}{2(u - \ell + k)} \int_{\ell}^{u} \pi(p) dp, \qquad \text{and} \qquad V_{nr}(\ell, u) = V_r(\ell, u) + \frac{k}{u - \ell + k} \pi(u).$$

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 $\overline{p}_C > \underline{p}_C > \underline{p}_U$ (which is 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*

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

These equilibria all result in identical price paths and payoffs.

(All proofs in Appendix A.) 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 most periods are to monopoly profits. However, the longer the cycle, the less frequently it has to be reset (which is costly as it requires a period of zero profits). In Appendix C we discuss how these results generalize. With less than infinitely elastic residual demand, cycling survives if k is large enough to make resetting the cycle more attractive than being permanently beaten on price by k. With more than two players, cycles can be sustained if and only if there is exactly one U-type agent.

[18] As $\delta \to 1$, Maskin and Tirole show joint profit must exceed 50% of monopoly in symmetric Edgeworth MPEs.

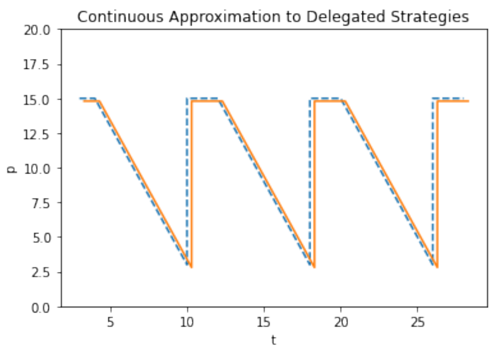

(a) Continuous Approximation

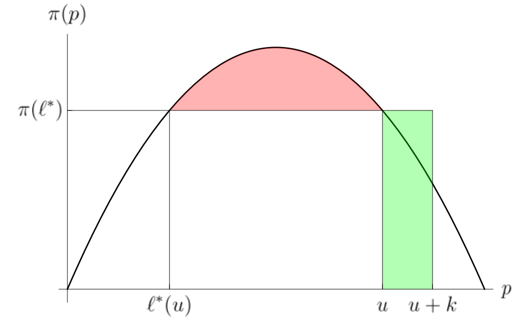

(b) The FOC for Choice of ℓ

Figure 13: The Continuous Approximation (left) & The FOC for Choice of ℓ (right).
Notes: The left figure shows the price path that our continuous approximation implies (see LHS of Figure 12 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 ℓ 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 ℓ, i.e., a longer cycle.

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:

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

These FOCs equate the marginal benefit of extending the cycle (LHS) to the marginal cost of extending the cycle (RHS). The tradeoff for the second FOC is illustrated in Figure 13 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 π 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 $\tilde{u}^* \leq u^*$.*

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. For an equilibrium between C and U, prices will spend at least half of the time *above* the monopoly price which, from a static perspective, is a Pareto loss to the economy: all merchants and consumers would be better off if prices were lowered. However, regularly exceeding the monopoly price is necessary for the tacitly collusive scheme between C and U in a dynamic setting.

Before we move on, we mention several limitations of our model. We assume the Buybox is infinitely price-elastic and restrict attention to duopoly settings. While our results are robust to finite price elasticities, cycling can be interrupted by additional entry as coordination on resetting becomes more difficult (see Appendix C). Furthermore, our model accounts for competition between products on Amazon (as opposed to between offers within products) only through the reduced-form demand curve $D(p)$. Thus, when we say that average prices are near monopoly, the monopoly price we refer to is that of a firm with total market power for a given product but possibly facing stiff competition from firms selling other products.

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[19] 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 B.

Plugging in the equilibrium parameter choices, the stage game of technology choice then has payoffs given by Table 4, 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 $V_f - c_U > V_{ff}$. 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 randomly choose prices. Note that this implies 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

[19]The assumption of random matching is meant to reflect two key features of the market: high turnover and the fact that sellers typically sell a large number of products. To the extent that there is high turnover (as illustrated in Figure 4a), sellers cannot predict their future competitors. Similarly, when you sell a large number of products, it is unlikely to be optimal to carefully select a strategy for each product. Instead, sellers are likely to use the same strategy for all products and only modify the parameters of this strategy. Indeed, repricing interfaces are carefully designed to allow for this kind of bulk strategy adaptation: e.g., many firms offer the possibility to ingest minimum and maximum prices for all products from a CSV file.

		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 4: 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 $V_f - c_U$.

population and 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 ϵ, 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 ϵ, the stationary distribution $\mu(\epsilon)$ will put a positive probability on all possible states. As $\epsilon \to 0$, KMR show that the stationary distribution converges to a unique limit μ^*. Returning to the game described in Table 4, KMR's Theorem 5 trivially implies:

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* $(0, \lceil \gamma N \rceil, N - \lceil \gamma N \rceil)$.

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, for any initial shares over the three repricing strategies, the population of strategies will converge to having, approximately, γN using U and $(1-\gamma)N$ using C. 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 (theoretically) merely pave the way for future price increases.

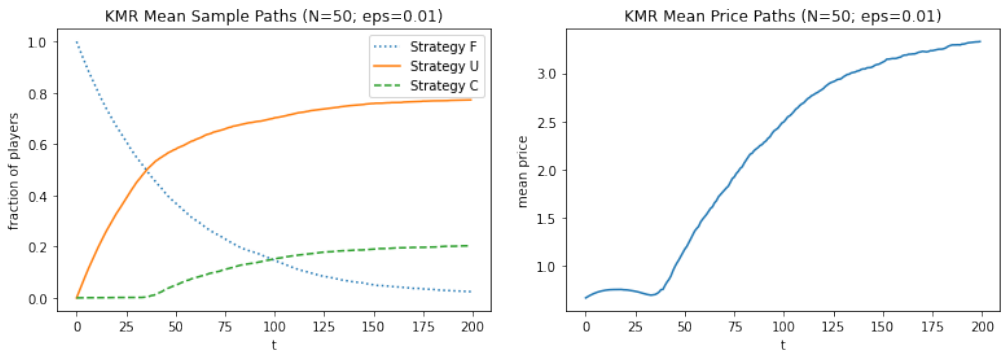

Figure 14: The Third-Party Marketplace Adjusts After Introduction of Repricers.
Notes: This figure illustrates the mean adoption paths of the three possible repricing strategies (left panel) and the resulting path of mean prices (right panel).

Simulation. To illustrate the evolution of the population of repricing agents, 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 14 depicts the resulting mean sample paths, which follow 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 cost and stay there. We soon see agents switching to cycling strategies to avoid earning zero profit. These strategies allow extraction of rents 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 whose prices are below the intended cycle. This inability ensures the market converges to a mix of undercutting and cycling strategies.

Finally, the example satisfies the conditions of Proposition 6 above: and indeed, in the right panel of Figure 14 the mean price dips briefly after initially increasing, only to shoot up eventually. While this rapid price increase is concerning, we emphasize that our simulation makes several simplifying assumptions that are unlikely to hold in reality. In particular, each product is sold by exactly two merchants. More realistically, a rapid price increase would soon be followed by a surge of entry into the market that could reverse the price increase. Furthermore, it is likely not in Amazon's interest to allow high prices on its platform, and as it controls the Buybox algorithm, it could potentially intervene to prevent prices from rising too much (Johnson et al., 2023). Thus, while the simulation is illustrative, it should not be taken as a prediction of the future evolution of the repricing economy.

Empirics. The above discussion suggests that repricing may initially increase welfare; however, if undercutting turns into cycling, prices may increase, and welfare decrease. So did cycling increase over time? We employ historical data from Keepa to answer this question. The advantage of the

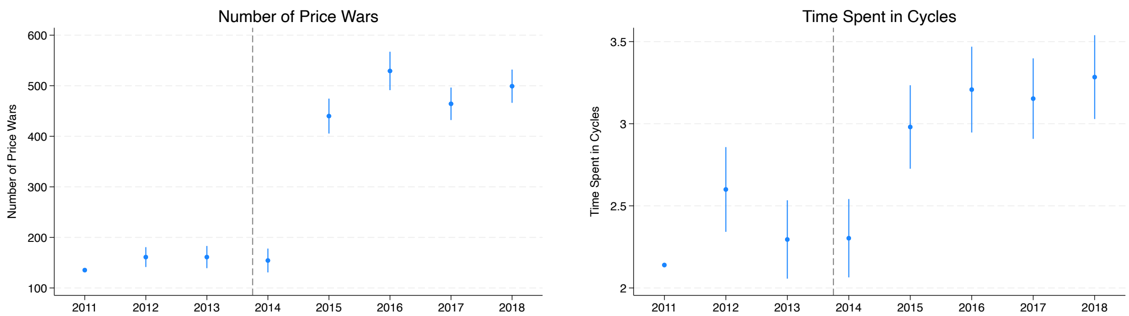

Figure 15: The Secular Increase in Price Wars and Cycling.
Notes: Each panel plots the coefficients and standard errors (clustered at the product level) from regressing a measure of cycling or price wars on product fixed-effects and year fixed-effects. Both price wars and cycling have increased since Amazon introduced the Subscription API in September 2013 (sellercentral.amazon.com).

Keepa data is that it extends back to 2011 (our proprietary data only extends back to 2018).

Our dataset consists of the 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 α_t is a set of year dummies and β_i is a set of product fixed effects. We report the year dummy coefficients in Figure 15. The dashed grey line marks the introduction of the Amazon MWS Subscription API, which allows subscription to price change notifications. We see that the number of price wars and time spent in cycles have increased over time after MWS API's introduction.

5 Conclusion

This paper employed unique high-frequency e-commerce data, a novel algorithmic cycle-recognition approach, and a model of equilibrium in delegated strategies to show 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 that would arise between undercutting merchants, some repricing companies have developed resetting strategies that regularly raise the price if it has fallen below some value. Thirdly, we found evidence that these (currently still relatively rare) resetting strategies effectively raise opponents' prices. Finally, resets happen much more frequently during the night when sales probabilities are lowest, indicating that repricers are aware of the strategic tradeoffs inherent to their resetting strategies. These conclusions are consistent with the marketing materials of repricing companies, which emphasize the importance of resetting in raising competitors' prices.

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 through which merchants have to enter their pricing rules allow *only* deterministic resets, effectively making it impossible to implement Maskin-Tirole strategies.

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. Indeed, currently cycling strategies are still rare, and hence they are very attractive because they perform well when undercutting strategies are widespread among sellers; however, if cycling strategies were to become more common, their performance would degrade as they would increasingly face other cyclers. Nevertheless, the market can support a stable mix of cycling and undercutting strategies in equilibrium as shown using an evolutionary model. Thus, the required miscoordination remains plausible as a long-run outcome. Furthermore, our model predicts that while introducing repricing strategies may initially lower prices, it can increase prices in the long run. Crucially, unlike Maskin and Tirole (1988), 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 the monopoly price.

Still, our model has several limitations. We assume there are only two merchants for each product. While cycles can be sustained with multiple competitors (see Appendix C), the conditions for this are quite restrictive. Thus, as collusive profits may induce entry, the model may overstate the extent to which cycling can be sustained. Furthermore, the model assumes agents match randomly. While consistent with evidence of high turnover on Amazon, merchants do in practice find themselves facing the same competitors regularly. Thus, there may be more scope for a firm to adapt to a competitor's strategy than our model suggests. Finally, we restrict our agents to choose the strategies currently observed in practice. However, more advanced strategies may be invented and adopted in the future.

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 competitors that rarely update their prices. Similarly, resetting algorithms are effective when a competitor is 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 prices are constrained in their response to cost and demand changes.

References

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.

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 (2024). Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. *Journal of Political Economy 132*(3), 723–771.

Banchio, M. and G. Mantegazza (2022). Adaptive Algorithms, Tacit Collusion, and Design for Competition. *arXiv*.

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

Brown, Z. Y. and A. MacKay (2023, May). Competition in Pricing Algorithms. *American Economic Journal: Microeconomics 15*(2), 109–56.

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.

Cooper, W. L., T. Homem-de Mello, and A. J. Kleywegt (2015). Learning and Pricing with Models That Do Not Explicitly Incorporate Competition. *Operations Research 63*(1), 86–103.

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.

Fershtman, C. and A. Fishman (1992). Price Cycles and Booms: Dynamic Search Equilibrium. *The American Economic Review 82*(5), 1221–1233.

Fershtman, C., K. L. Judd, and E. Kalai (1991). Observable Contracts: Strategic Delegation and Cooperation. *International Economic Review 32*(3), 551.

Freyaldenhoven, S., C. Hansen, J. Pérez Pérez, and J. M. Shapiro (2021, August). Visualization, Identification, and Estimation in the Linear Panel Event-Study Design. Working Paper 29170, National Bureau of Economic Research.

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. (2022). The Effect of Outsourcing Pricing Algorithms on Market Competition. *Management Science 68*(9), 6889–6906.

Johnson, J. P., A. Rhodes, and M. Wildenbeest (2023). Platform Design When Sellers Use Pricing Algorithms. *Econometrica 91*(5), 1841–1879.

Jungle Scout (2023). State of the amazon seller 2022. https://www.junglescout.com/amazon-seller-report/2022-results/. Accessed: 2024-01-14.

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 (2023). Entry into Two-Sided Markets Shaped By Platform-Guided Search. *Working Paper*.

Leisten, M. (2022, 2). 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.

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, Nov). 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.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199. Themed Issue: Treatment Effect 1.

Tappata, M. (2009). Rockets and feathers: Understanding asymmetric pricing. *The RAND Journal of Economics 40*(4), 673–687.

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.

Yu, C. (2023, October). The Welfare Effects of Sponsored Product Advertising. Accessed: 2024-01-14.

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

Online Appendix

A 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*

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

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 ϵ. 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 ℓ^*. □

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 π. Secondly, the determinant of the Hessian of $V_r(\cdot, \cdot)$ satisfies

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

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 π is single-peaked, $\ell^* < p_m < u^*$ and hence $D > 0$. □

Proposition 3. *If π 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 ℓ must increase or u decrease, and as $\tilde{\pi}(\cdot)$ will also be symmetric around the same p_m we thus have that both ℓ 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 π around p_m 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 γ 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. $\square$

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 α_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. $\square$

B 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

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

$\square$

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 ϵ. 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 ϵ. 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:

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

□

C Model Extensions

Less Than Perfectly Elastic Demand. While the main text assumes $D_i(p_i, p_j) = 1\{p_i < p_j\}D(p_i) + 1\{p_i = p_j\}D(p_i)/2$ for $i = 1, 2$ and $j \neq i$, we now show that the intuition behind our key result (that there is an equilibrium between C and U agents that results in cycling near monopoly price) does not rely on a knife-edge assumption of infinitely elastic demand. Instead, as long as demand is sufficiently elastic relative to the undercutting step size k, the main intuition continues to hold in a setting with non-homogeneous products, i.e., where $D_i(p_i, p_j) = A - bp_i + ap_j$. When residual demand becomes too inelastic, the cost of resetting the cycle becomes large relative to the loss from being permanently beaten on price by your opponent: in particular, for any fixed (a, b), as $k \to 0$, the costs of your opponent beating you on price disappear. Thus, with sufficiently small undercutting steps, the C agent will deviate from a cycling equilibrium by simply choosing to price at (or near) the monopoly price level[20]; the U agent (with unchanged strategy) will then (in the limit) match C's price.

Let $D(p) = D_i(p, p) + D_j(p, p)$. In the limit as $k \to 0$, note that in a cycling equilibrium the C agent has payoff

$$V_r(\ell, u) = \frac{1}{u - \ell + \Delta} \int_\ell^u (p - c)D(p)/2dp < (p^m - c)D(p^m)/2.$$

The inequality holds because (i) the resetting party loses a period Δ during which they reset the cycle[21] and (ii) there are inherent losses from price drifting away from the optimal price (which in turn exist in equilibrium because the resetting party trades them of against (i)).

For given $k > 0$, our simplifying assumption of a continuous price decrease during the declining part of any possible cycle ceases to be true, and we hence lose the ability to tackle the setting with the analytical machinery built up in the main text. However, we can solve the game numerically by forward simulating price paths under certain strategies. For $D_i(p_i, p_j) = 40 - 8p_i + 6p_j$, we find that with $k > 0.17$, a cycling equilibrium exists. For $k < 0.17$, the equilibrium cannot be sustained as the deviation discussed above becomes too attractive. We note that for these parameters, an opponent undercutting your price by 0.17 as opposed to matching your price leads to an $\approx 5.3\%$ decline in your demand. By contrast, Figure 2 shows that for the Buybox, even a 1% price cut by your opponent would lower your demand by more than 50%. Hence, the empirical setting seems comfortably in the regime in which theory predicts a cycling equilibrium.

More Than Two Players. When there are multiple players, its natural to have agents react to the current minimum price (as opposed to, say, a particular competitor's price.) However, if there are multiple U agents, each trying to undercut the currently lowest price, then no cycle can occur because the U agents would keep fixating on each other's prices even as a C agent attempts to reset the cycle.

If there is exactly one U agent, simulations confirm that cycles can still occur. In particular, with $n = 3$ (with types C, C, U), and the assumptions of the simulation in the main text (zero marginal cost, $D(p) = 20 - p$, $k = 1$ and all prices and parameter choices being restricted to a grid $\mathbb{N} \cap [0, 20]$), we find that a cycle between $p = 0$ and $p = 17$ is an (but not the unique) equilibrium.

The cycle has extended relative to the discussion in the main text. This is because the strategic incentives have shifted: with multiple C agents, deviating to a lower $\underline{p}_C$ involves not only an extension of the cycle (for U will only come up to reset the cycle once all C agents have increased their prices)

[20]Recall that C can implement a fixed price p by setting $\overline{p} = p$ and $\underline{p} > \overline{p}$.

[21]Note that we call this period k in the main text as it plays an analogous role to k, but here this notational overloading would be confusing.

but also an attractive period during which profits only need to be shared with U (as opposed to all other players.) To counteract this incentive to extend the cycle downwards, prices at the bottom of the cycle need to be anchored closer to marginal cost than previously.

Prices at the top of the cycle face no such constraint: supposing both C agents set the same $\underline{p}_C$, the effective top of the cycle will be the minimum of their reset prices. Hence, we are faced with an embarrasment of riches when it comes to equilibria: if my opponent sets their reset price to be x, then my payoffs are independent of my reset price as long as it stays weakly above x, and hence I would be willing to set my reset price to x in turn. This multiplicity of equilibria, however, is easily resolved by assuming that agents break ties in favor of reset prices that would lead to higher profits if they were implemented by their opponents, which pushes in the direction of longer cycles.

D Cycle & Price War Recognition

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*[22] *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.

[22] The 'reaction time' is the amount of time between a competitor's price change and my own (next) price change; the 'average' is the average of these reaction times since the start of the cycle.

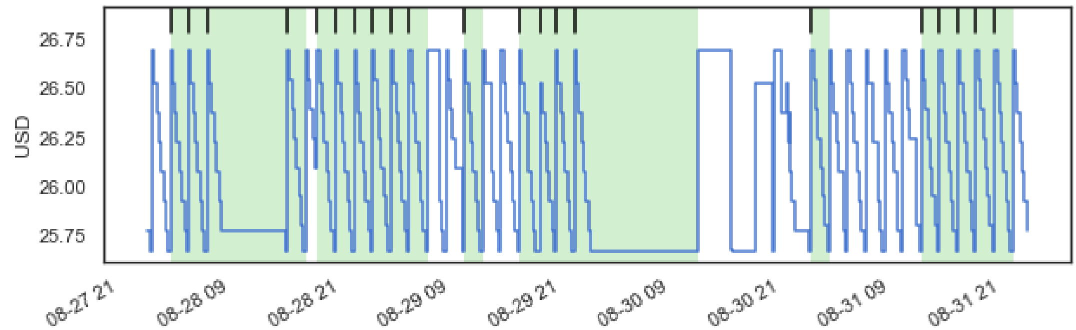

Figure D.1: Output of The Cycle Recognition Algorithm.
Notes: This figure illustrates the output of the cycle recognition algorithm using an example price path. The periods shaded in green are declared to be ones in which the offer is cycling by the algorithm.

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, 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.

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.*

E Robustness of Event Studies

We report various figures that investigate the robustness of our results to alternative specifications and estimators.

Repricer Effects on Sales and Profits. For a subset of 4,035 offers, we observe the number of sales and the reported unit costs (and hence profits) associated with these sales. We separately explore the effects of repricer activation (including on sales and profits) for these offers in Figure E.5. We caution that the experience of these offers is quite distinct from most offers. While we still see evidence suggesting a price decline, the effect is much smaller (and not statistically significant) for these offers. More intriguingly, the effect on Buybox share is much smaller and not statistically significant. Importantly, the Buybox results seem to indicate that relative to our understanding of the data, there may be an additional one-day lag in our observation of repricing status; this is inconsistent with our understanding of the DGP, but we may be missing a caching process that delays the observation of repricing status. It is clear from the plots that such caching would not change our conclusion on any other outcome (or indeed any outcome for the full sample), but it would change our conclusion on Buybox share for the subset of offers for which we observe sales and profits. In particular, it would indicate that Buybox share declines after repricer activation. Consistent with this, we find a decline in sales and profits in the final two panels of Figure E.5. Overall, we caution that the offers for which we observe sales and profits are not representative of the entire sample: reporting cost information to the repricing company indicates a high degree of sophistication and may also indicate that the merchant is using distinct pricing strategies that rely more on cost information (e.g., passing along cost information may be more valuable for merchants without competitors.)

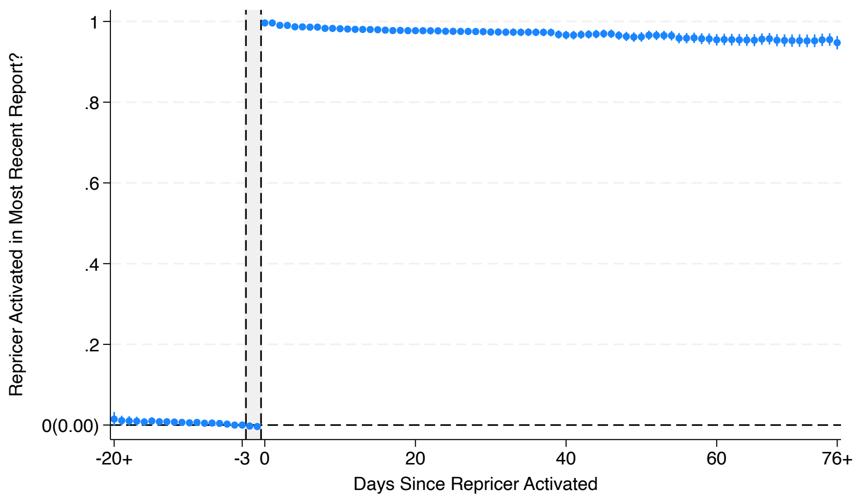

Figure E.1: Stickyness of Repricer Activation.

Notes: This figure exhibits results from an event study that measures the effects of initial repricer activation on repricer activation itself. The horizontal axis measures days since the repricer was activated (with 0 being the day of activation.) Due to sampling frequency limitations, repricing status is unknown during the time window indicated by light gray shading; thus, we normalize the coefficient that measures the treatment effect three days before (certain) treatment to zero to avoid expressing treatment effects relative to a period where treatment has already started for some units. The vertical axis measures the effect of activating the repricer on the outcome variable of interest, with a zero value indicating no effect; we also provide the mean of the outcome variable three days before treatment in the parenthetical label. The specification used is identical to that underlying Figure 7 and Table 2.

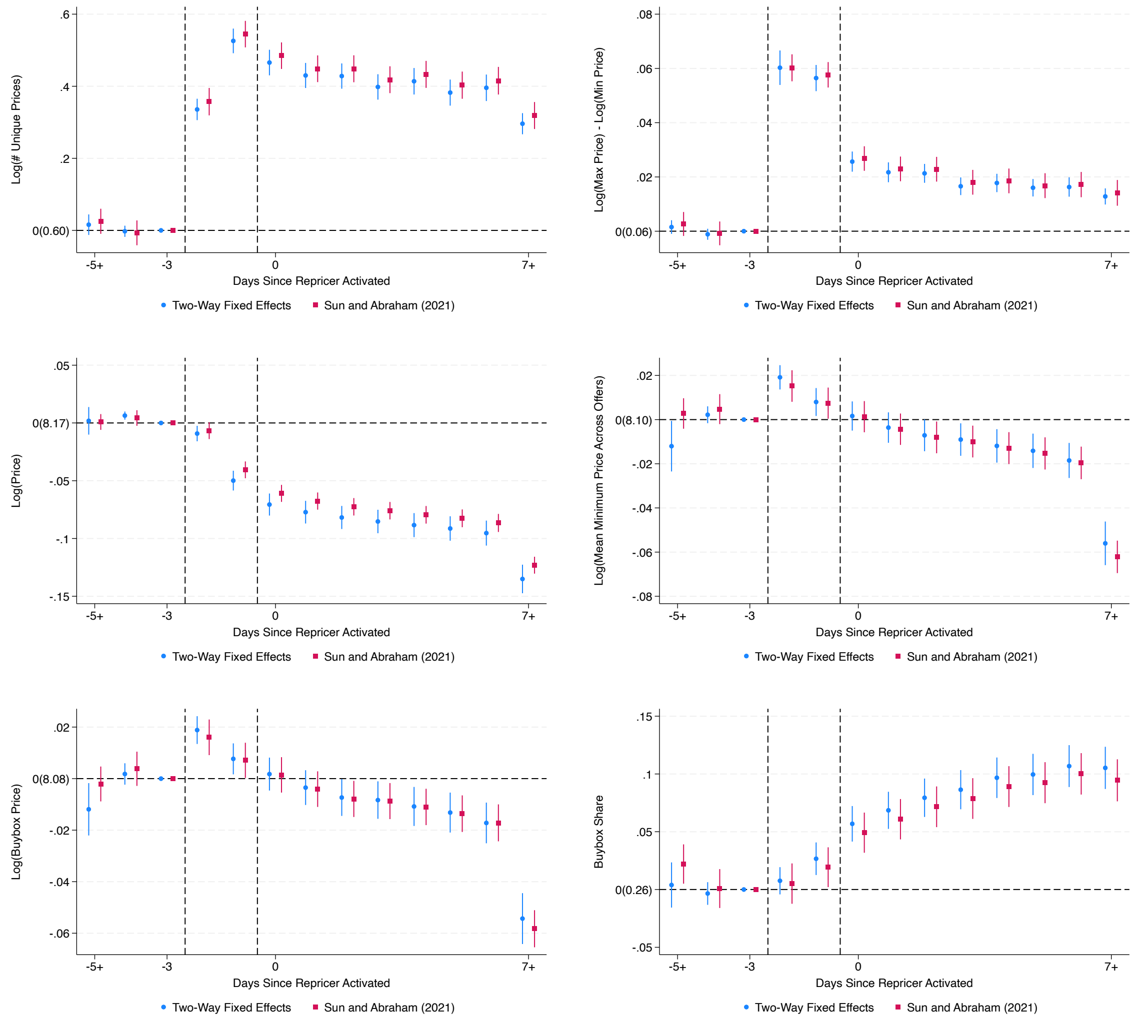

Figure E.2: Repricer Activation Event Study Is Robust To Alternative Estimators.

Notes: This figure repeats the results of the activation event study figure in the main text but additionally implements the Sun and Abraham (2021) estimator to investigate robustness of our results to treatment effect heterogeneity. Due to the considerable computational requirements of the procedure (which requires introducing additional regressors on the scale of # cohorts x # lags/leads), we are forced to limit the number of dynamic treatment effects we consider in this figure.

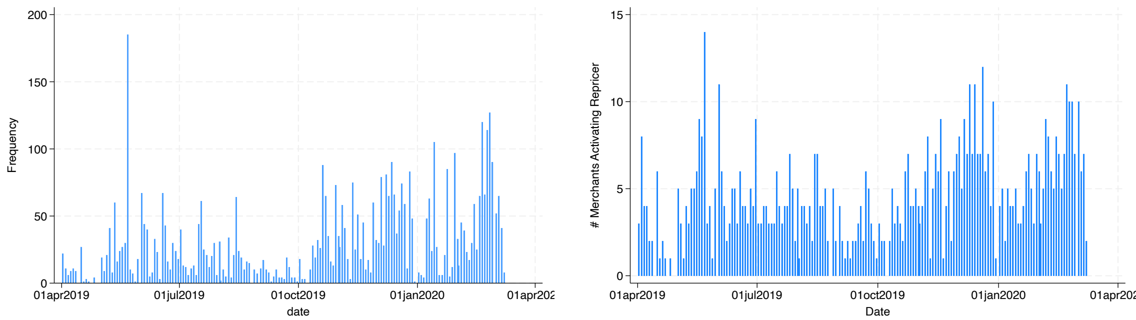

Figure E.3: Distribution of Repricer Activation Events
Notes: This figure illustrates distribution of repricer activation events used in the estimation of the repricer activation event study. The left panel is at the offer-level, and the right panel is at the source merchant level.

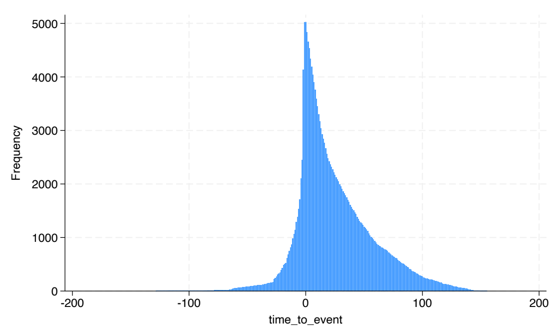

Figure E.4: Distribution of Time To Repricer Activation.
Notes: This figure illustrates distribution of the number of days until repricer activation in the panel used to estimate the repricer activation event study. We see that there is much more mass to the right of zero, indicating that we observe offers for a long time after the repricer was activated but not for a long time before. For this reason, we choose to limit the number of leads of the treatment effect to be less than the number of lags when we estimate our main specification.

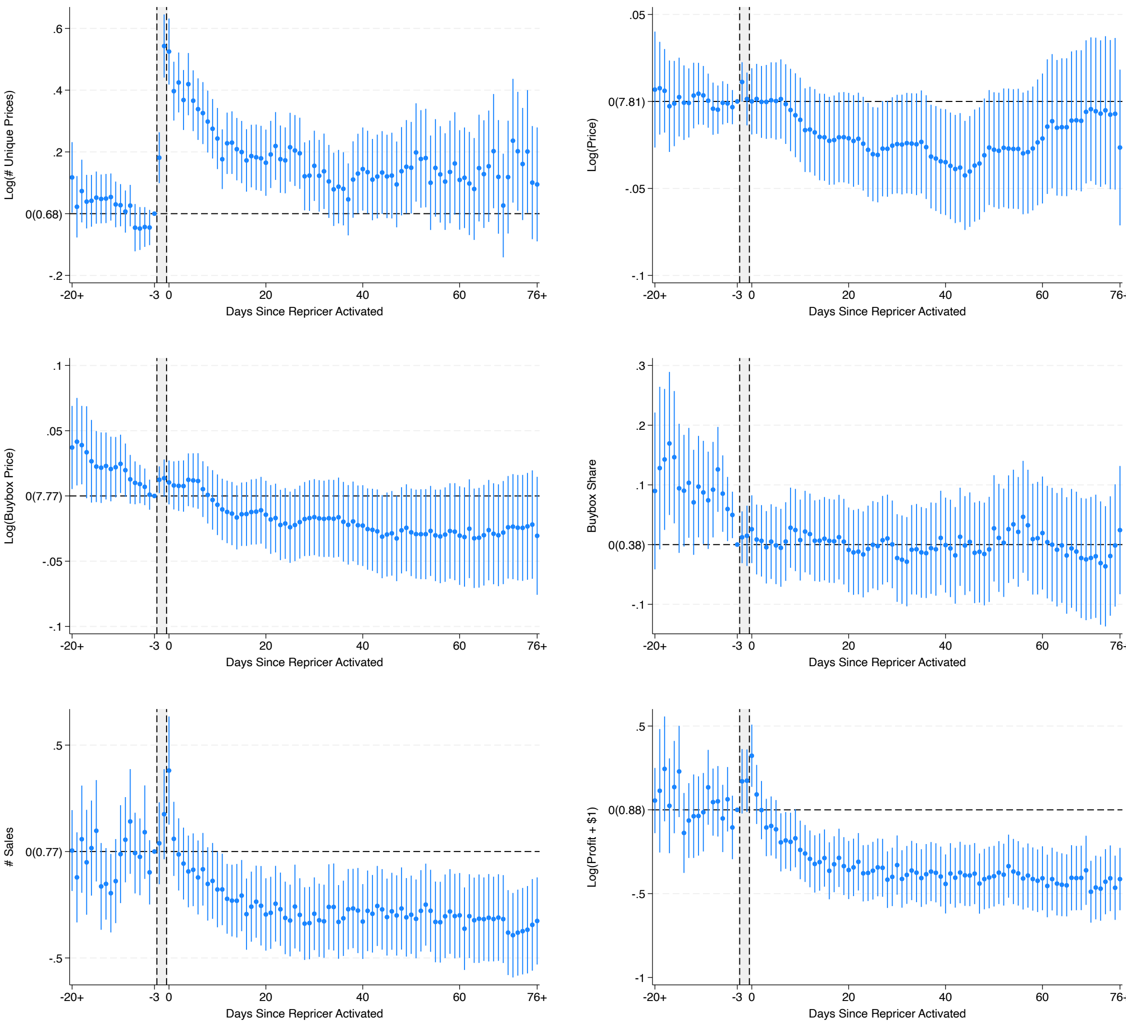

Figure E.5: Repricer Activation on Offers Reporting Sales And Cost Data.

Notes: This figure exhibits results from an event study that measures the effects of initial repricer activation; it differs from 7 only in that the sample has been limited to offers reporting cost and sales information (4,035 offers report this information.) We separate out these offers because they behave quite differently from the remaining offers. The horizontal axis measures days since the repricer was activated (with 0 being the day of activation.) Due to sampling frequency limitations, repricing status is unknown during the time window indicated by light gray shading; thus, we normalize the coefficient that measures the treatment effect three days before (certain) treatment to zero to avoid expressing treatment effects relative to a period where treatment has already started for some units. The vertical axis measures the effect of activating the repricer on the outcome variable of interest, with a zero value indicating no effect; we also provide the mean of the outcome variable three days before treatment in the parenthetical label. Each blue dot corresponds to a coefficient β_s in (1), and the bars indicate a 95% confidence interval derived from standard errors clustered at the offer-level.

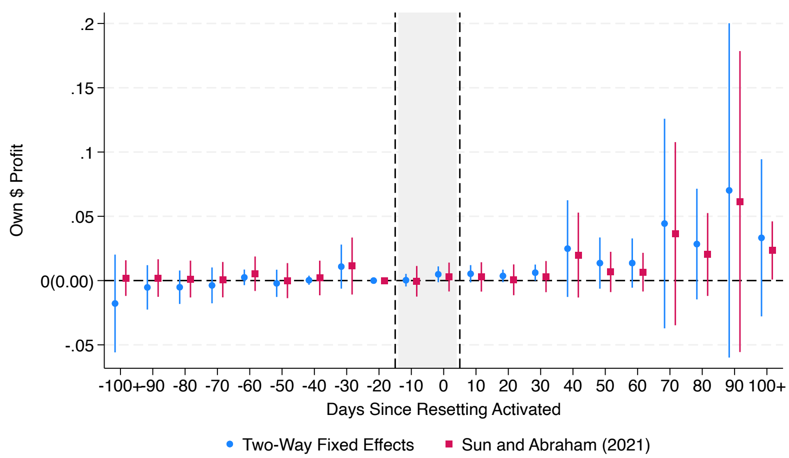

Figure E.6: Effect of Resetting on Profits.

Notes: This figure exhibits results from an event study that measures the effects of initial repricer activation on profits (measured in USD/day). We note that these results are based on the (possibly selected) sample of offers for which merchants report cost information; profit is calculated as quantity shipped times price minus seller-reported cost.

F Repricing Strategy Availability

The main text introduces a model in which merchants can choose between an undercutting strategy (which simply undercuts the opponent's price by a fixed amount) and a resetting strategy (which does this but also resets the cycle when the opponent's price is too low.) We now provide evidence that these strategies are indeed available to merchants, and that they are by far the most common strategies offered by repricing companies. To this end, we provide screenshots from the interfaces of all repricers listed in JungleScout's list of the "Best Amazon Repricer Tools for 2022" (junglescout.com) plus two repricers that were more common in the sample period.

Before we proceed, we summarize our findings. Firstly, there is a surprising amount of uniformity in what strategies repricers offer to their clients. Secondly, all repricers offer U-type strategies, and most offer C-type strategies. Thirdly, one strategy consistently offered that we do not model is a variant of C in which cycles are reset not when a certain price is reached, but instead at a regular interval; this strategy clearly yields qualitatively similar results as those in the main text (including prices near monopoly levels) but offers sellers less control over the length of the cycle in exchange for a less costly resetting period. Finally, another strategy easily implementable but not covered by our model is a matching strategy, in which the merchant simply matches the price of the competitor. This strategy is not discussed in our model because it does not generate easily detectable pricing patterns; indeed, if two merchants match each other's prices, they may quickly converge to pricing at the monopoly level, and the resulting stable prices would be hard to distinguish from Nash-Bertrand prices without more reliable data on margins.

We now go through the repricing interfaces one by one, beginning with a detailed discussion of the Informed.com repricing interface presented in Figure F.1. The interface begins by letting a merchant choose the type of offer you want to compete with (FBA, FBM, or all.) Then, he can choose how to price relative to your competition: he can price below, price above or match; if he prices below or above, he needs to specify by how much either in units of currency or as a percentage of the current price. He can also choose to set these settings differently for FBA, SFP or Amazon competitors, and specify what to do if there is no competition.

So far, the interface thus allows the merchant to implement a U-type strategy as discussed in the main text: he simply chooses to compete against all, and undercut by a fixed amount. However, the interface also allows him to implement a C-type strategy: he can choose what to do if his competition is below his minimum price or matches it. If he sets these fields to 'Use max price', he is effectively implementing a C-type strategy: as soon as he reaches the bottom of the cycle, he will set his price back up to his maximum price.

However, the interface does not stop here. He can also differently price if he is out of stock, or prevent the repricer from lowering his price when he is already winning the Buybox. Finally, he is given fine-grained control over who the repricer considers a competitor: for instance, backordered competitors may not be eligible for the Buybox, and hence he may choose to ignore them completely.

To summarize, the Informed.com interface allows implementation of U-type and C-type strategies,

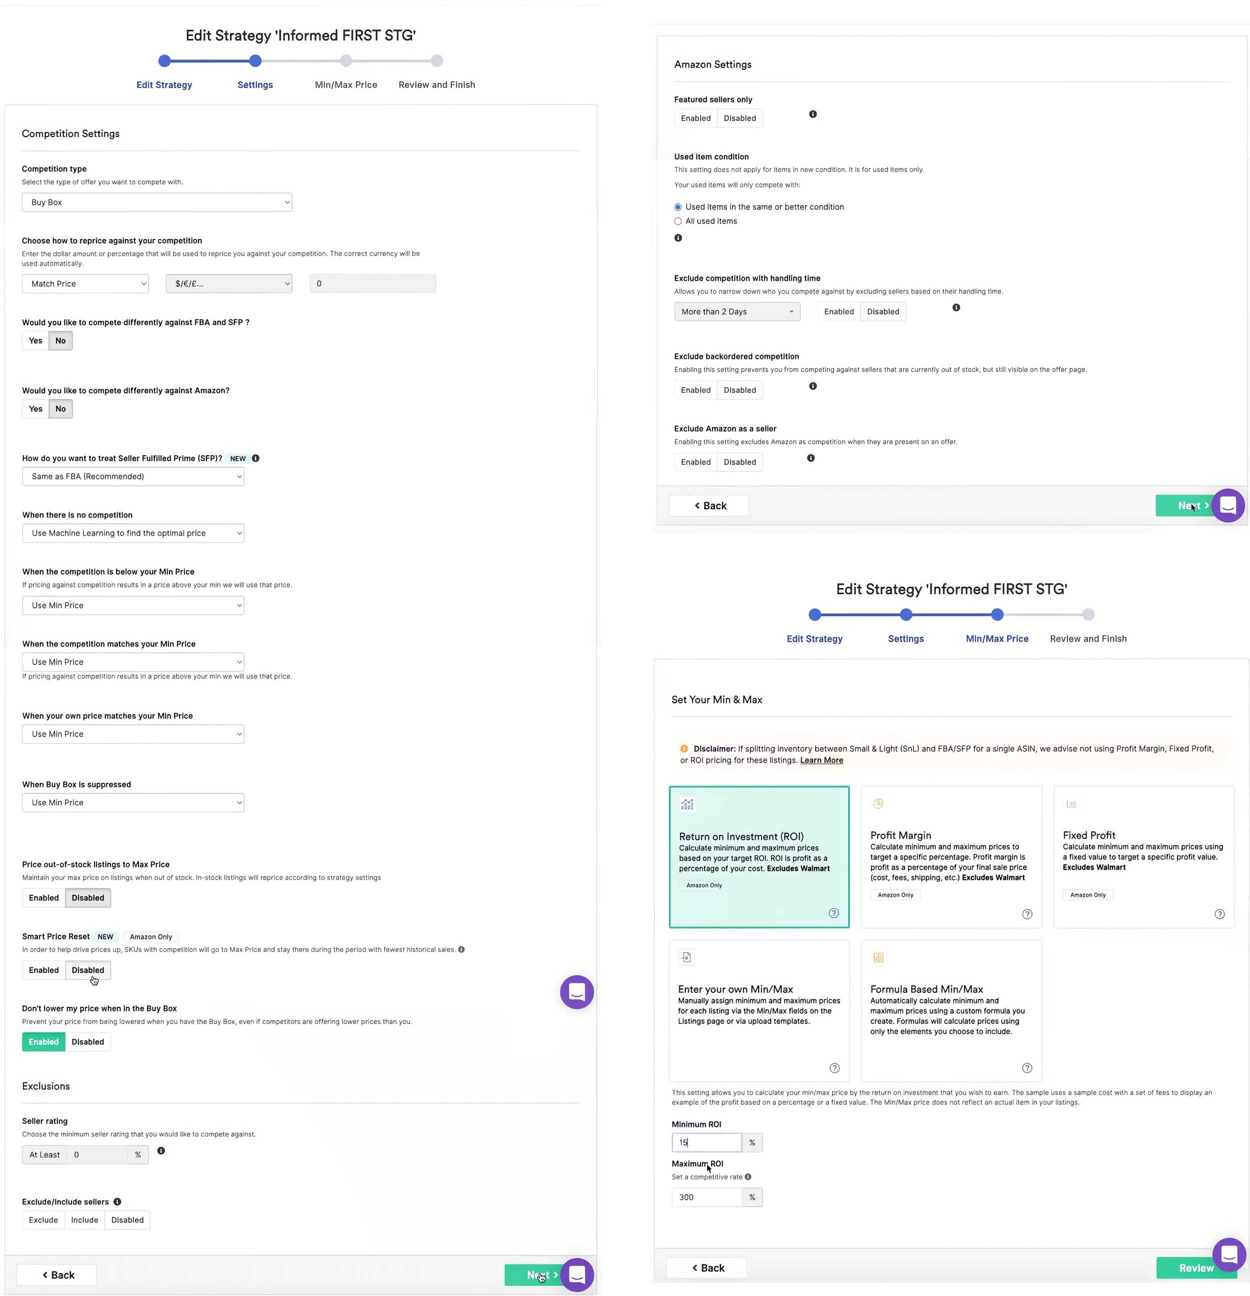

Figure F.1: Repricing Interface: Informed.

Notes: This figure depicts the repricing interface for custom strategy creation at Informed. The interface (see the left panel) explicitly offers a 'Smart Price Reset' that automatically resets the price to the maximum during the period with the fewest historical sales 'in order to help drive prices up.'

but it also contains some additional levers that the model abstracts away from. Still, these levers would not allow for the implementation of Edgeworth-style strategies (which require randomizations and jumps), and they mostly serve to deal with situations where there is multiple offers and only some should be considered competitors.

Moving on to the Aura.com repricing interface in Figures F.2 and F.3, we see that it is very similar to the Informed.com interface. The merchant begins by defining who his competitors are in the 'Competitors' section – for instance, he may not want to compete with Amazon itself. Next, the merchant can choose to price below, match, or price above his competition, and by how much. Again, options are given for setting a different pricing strategy depending on the competitors' fulfillment method. Most importantly, C-type strategies are again implementable by setting 'When the competition is below your min price' to 'Use max price.' Finally, the maintenance section in Figure F.3 introduces a new complication: merchants can choose to raise their price while in the Buybox, presumably to find the maximum price at which they still emerge as Buybox winner. This would never be a useful strategy against a type U or C strategy. Instead, raising your price after acquiring the Buybox is useful when your opponent is playing a fixed price, and you do not know by how much to undercut them. The model assumes that k (the minimum unit of currency) is sufficiently large that you always win the Buybox when you undercut by k, but a \$0.01 undercut may not be enough to win the Buybox if (e.g.) your competitor has more reviews. For these cases, automated discovery of the optimal undercutting amount is useful. Still, this complication does not affect the main results of the model; indeed we show in Appendix C that the model's main results are robust to the Buybox having less than infinite price elasticity.

To summarize, the Aura.com interface again allows implementation of U-type and C-type strategies, and it again has some additional levers which are mostly irrelevant to our model: selecting your competition (irrelevant when you have one competitor) and strategies to deal with the need to discover the optimal undercut (assumed away in the model.)

In addition to the repricing interfaces of Informed and Aura, we also exhibit the interfaces of RepricerExpress (in Figure 5) and BQool (in Figures F.4 and F.5). These interfaces are again very similar to those of Informed and Aura, and they allow the implementation of U-type and C-type strategies. However, they do allow for one more strategy which we observe in the data but do not explicitly model: resets at regular times. At RepricerExpress this strategy is called 'Sleep Mode' (see bottom right of Figure 5) and BQool calls this a 'Repeated Schedule' (see right of Figure F.5). Indeed, we did not discuss it above, but the Informed interface has a slightly more advanced version of this functionality: when merchants enable 'Smart Price Reset', Informed will automatically reset the price to the maximum during the period with the fewest historical sales 'in order to help drive prices up.' Similarly, SellerSnap's Yo-Yo Repricing Rule (Figure F.8) regularly raises the price to the maximum before starting repricing again.

Finally, we need to discuss Repricer.com (Figure F.6) and SellerSnap (Figure F.7). Repricer.com clearly allows for implementation of U-type strategies, but cannot implement C-type strategies. SellerSnap is slightly more opaque and offers an 'AI Repricer' together with various strategies that in practice amount to a U-type strategy or perhaps a price matching strategy. Most interestingly, as illustrated in Figure F.8, SellerSnap offers a 'Yo-Yo Repricing Rule' which is a C-type strategy that regularly raises prices – this is an example of the kind of regular reset strategy that intuitively leads to similar outcomes as a price-based resetting strategy.

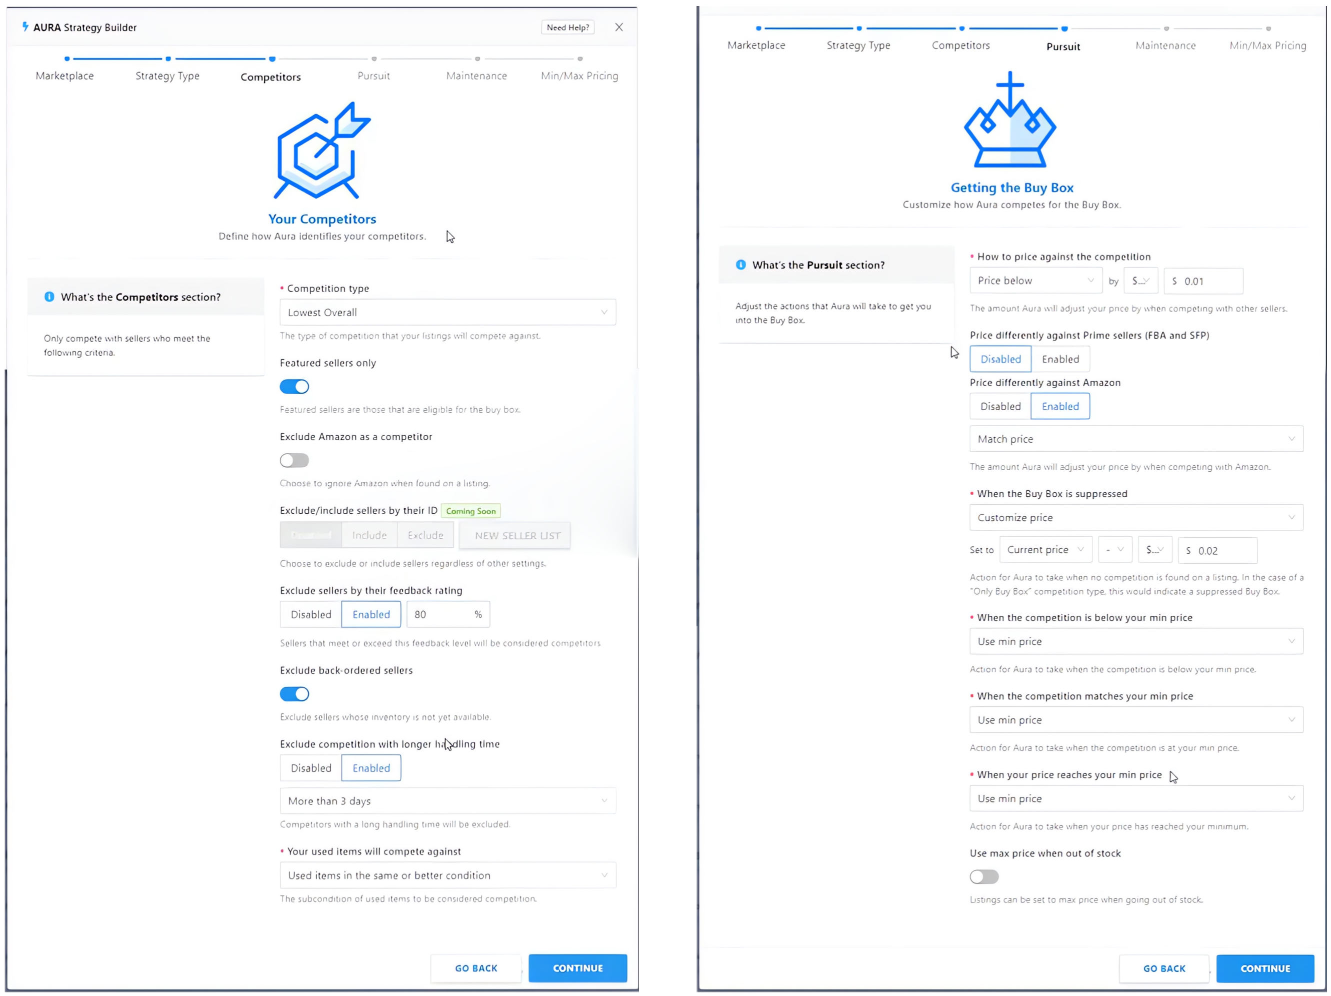

Figure F.2: Repricing Interface: Aura (Part I).

Notes: This figure depicts the repricing interface for custom strategy creation at Aura.

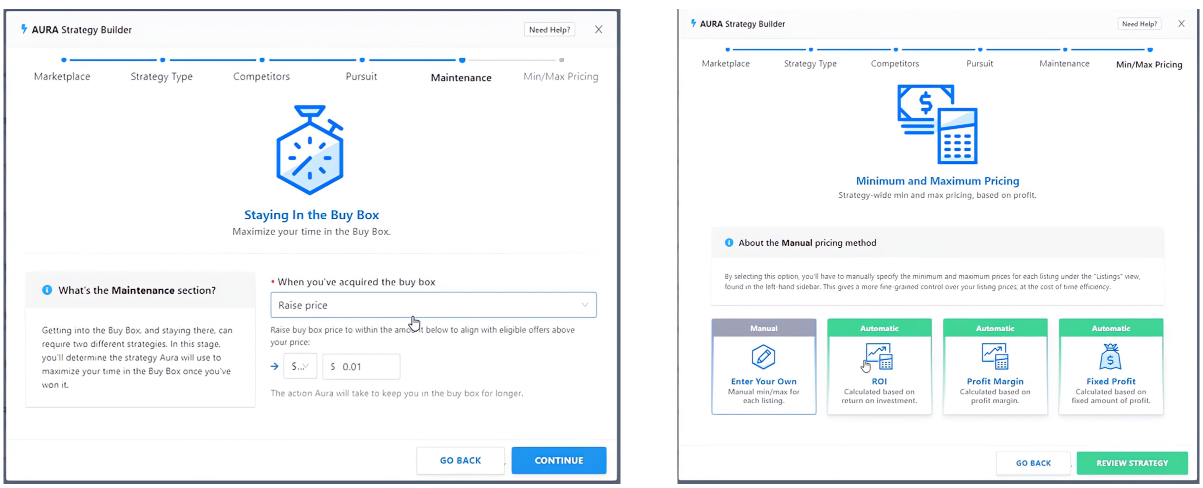

Figure F.3: Repricing Interface: Aura (Part II).
Notes: This figure depicts the repricing interface for custom strategy creation at Aura.

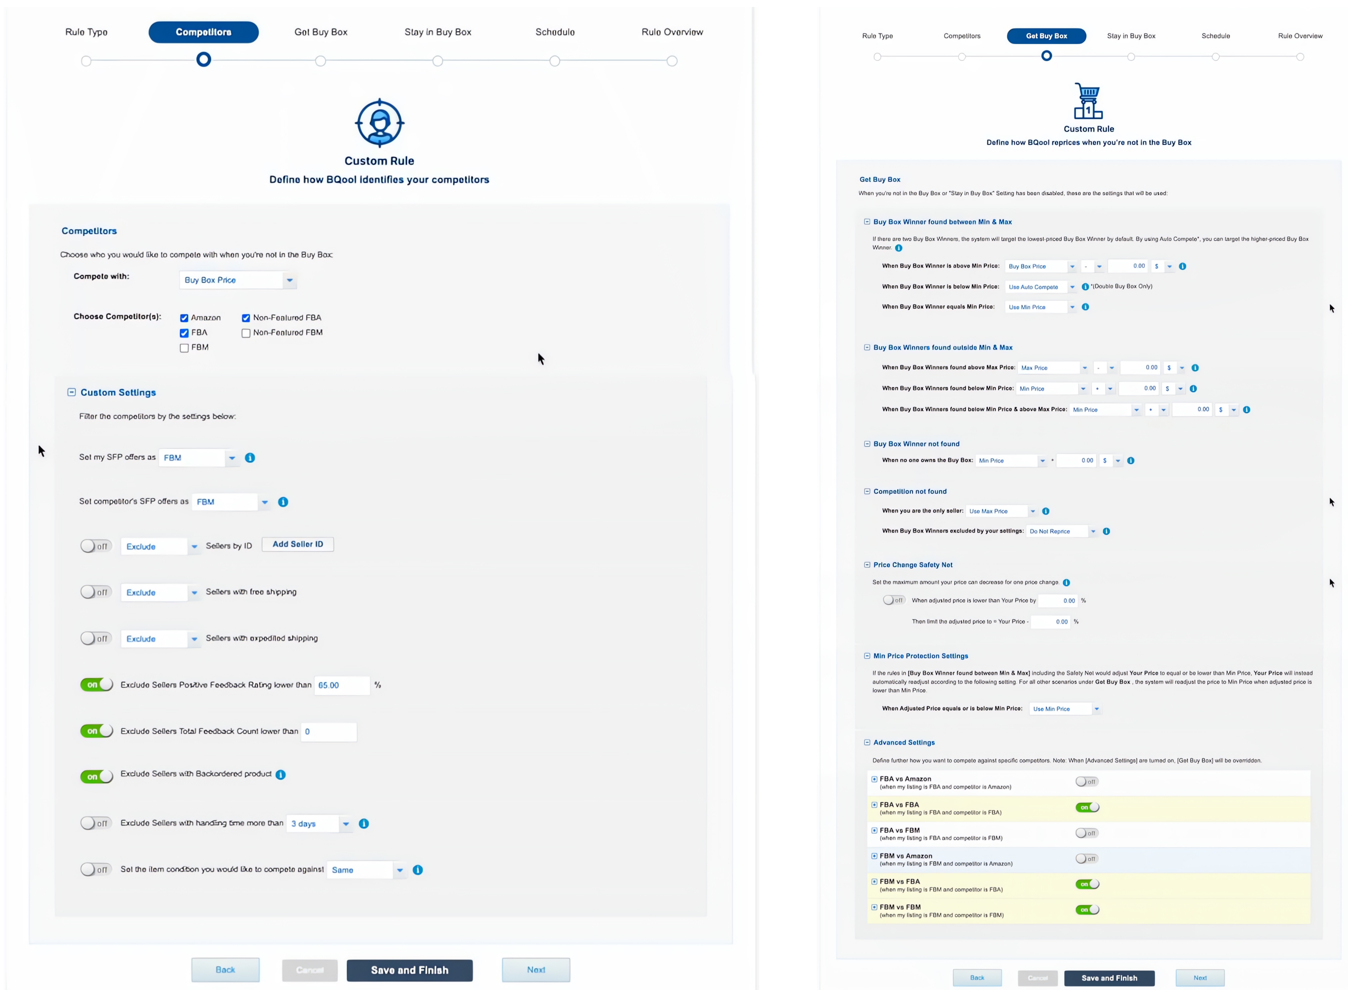

Figure F.4: Repricing Interface: BQool (Part I).

Notes: This figure depicts the repricing interface for custom strategy creation at BQool.

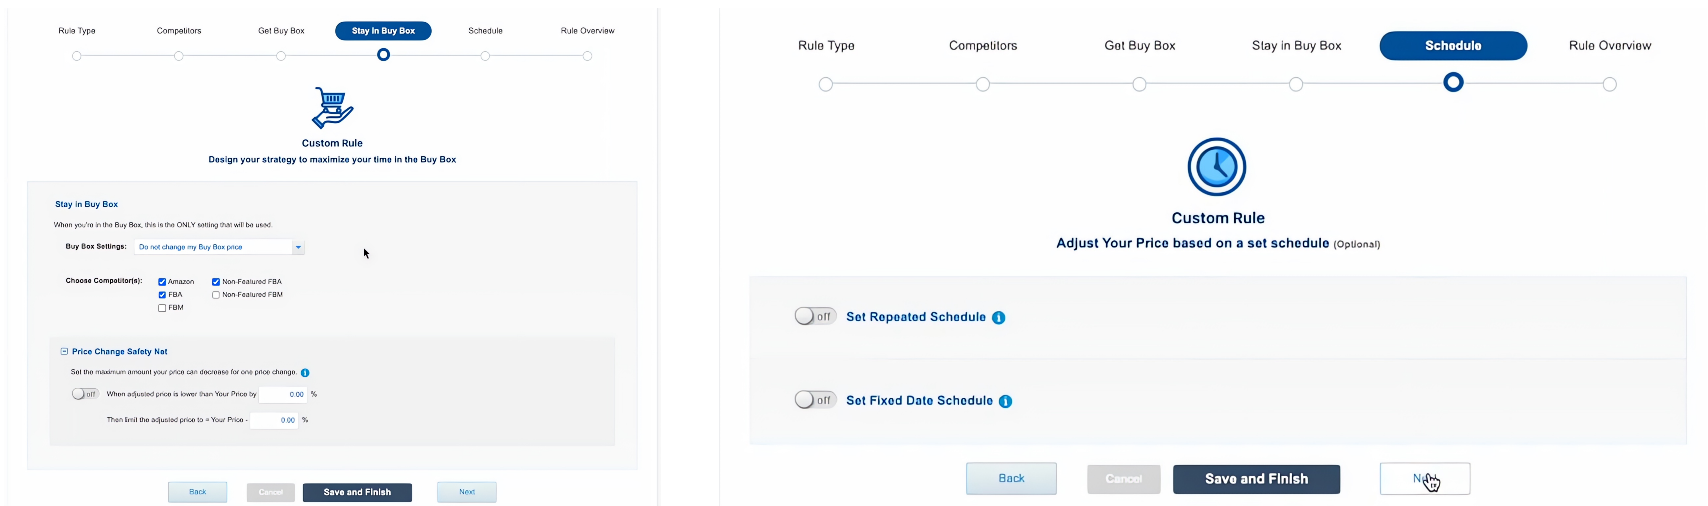

Figure F.5: Repricing Interface: BQool (Part II).

Notes: This figure depicts the repricing interface for custom strategy creation at BQool.

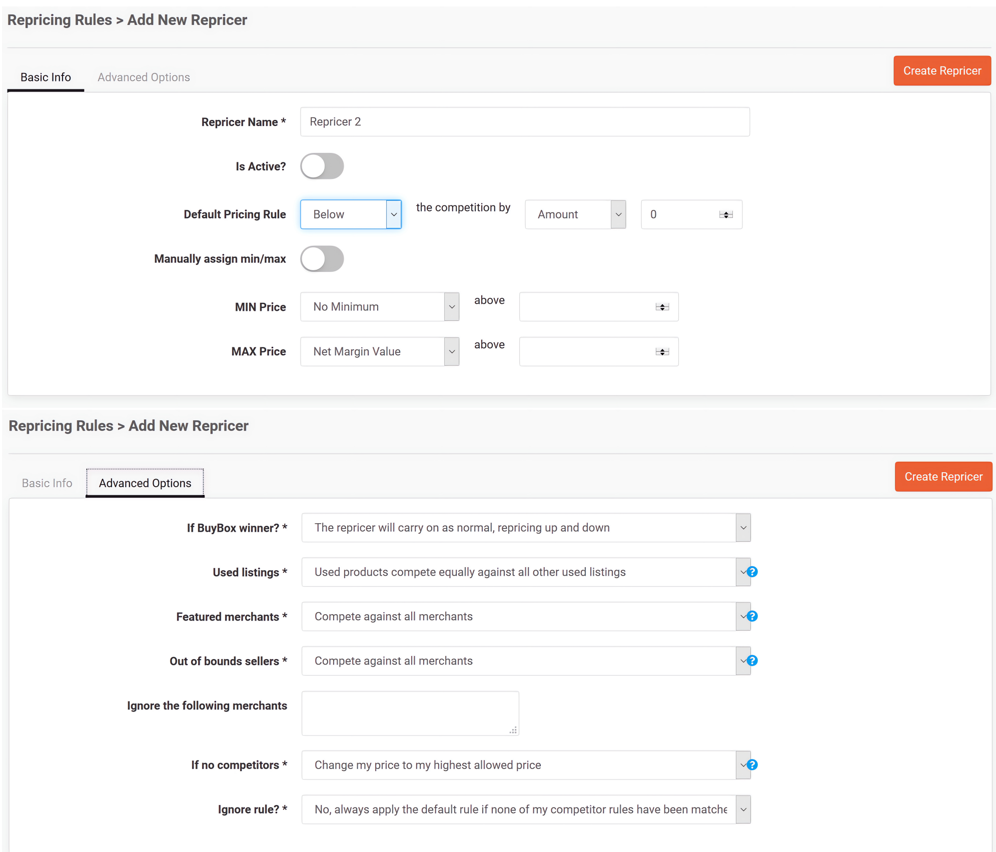

Figure F.6: Repricing Interface: Repricer.com.

Notes: This figure depicts the repricing interface for custom strategy creation at repricer.com.

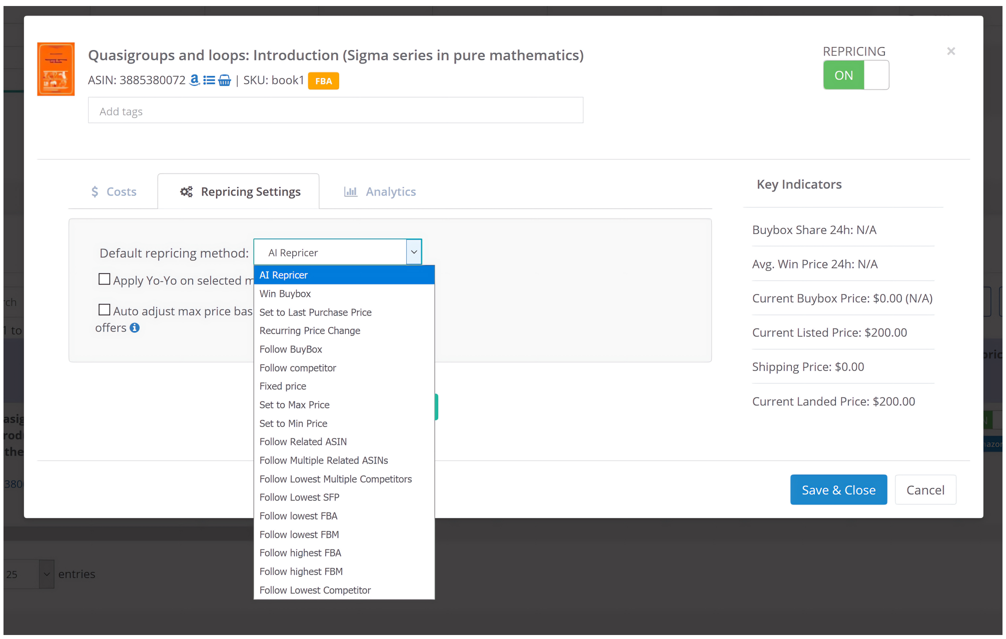

Figure F.7: Repricing Interface: SellerSnap.
Notes: This figure depicts the repricing interface for custom strategy creation at sellersnap.com.

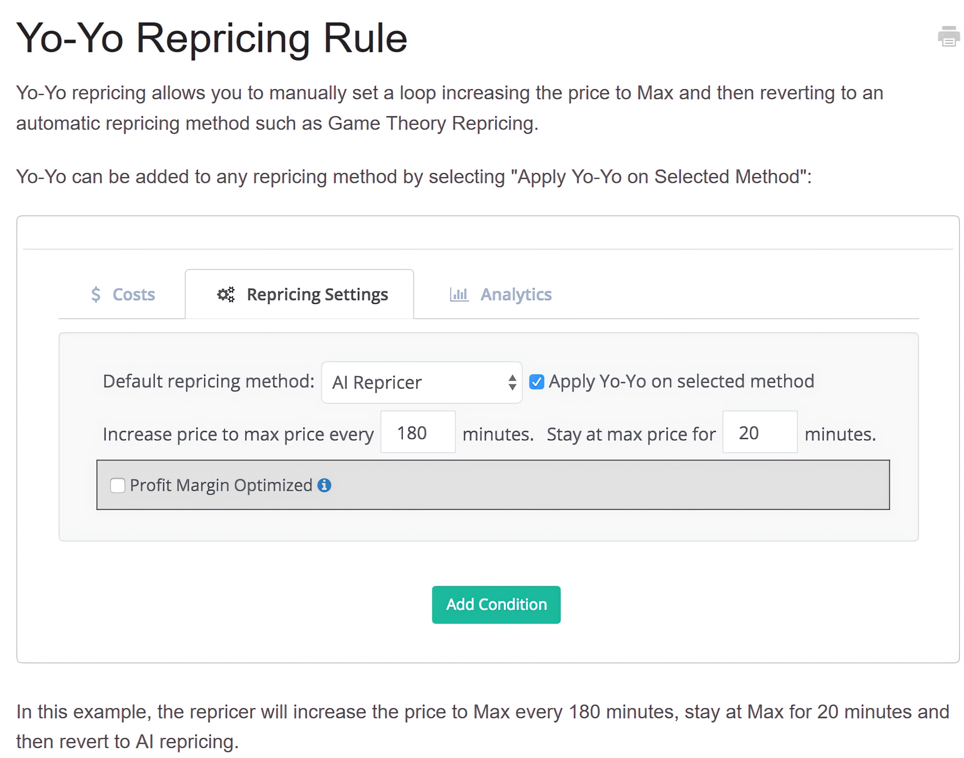

Figure F.8: Explanation of Yo-Yo Repricing Rule.

Notes: This figure depicts the 'Yo-Yo' repricing strategy at sellersnap.com. This strategy allows the merchant to regularly raise prices to the maximum at a pre-specified period. For instance, the depicted configuration raises prices to maximum for 20 minutes every 180 minutes.

Answered Mar 17, 2018

1. **Never use amazons own repricer**.

2. **Identify trends and adapt**.

   - **Can you raise the price during the night?** Colluding with other repricers for night bumps is very profitable.

   - Listing only has **1 competitor AND competitor has low stock?** Let them win the buy box and raise the price when they run out.

   - **What is your target price?** This should be from equivalent, competing products. Why? It doesn't matter if you win the buybox when your listing is so expensive that amazon won't show it in search results anyway.

   - **Who is your competition within a listing?** Should you be competing with Amazon? Should you be competing with other FBA inventory? Should you compete with self-fulfilled merchandise?

   - **How fast does competition reprice and is it a race to the bottom?** If you reprice faster AND competition will race to their minimum then it's more profitable to alternate between being 1 cent above them and 1 cent below them. The majority of the time we will have the BB and the price won't drop dramatically for everybody.

- Are you using the right technology? The best Amazon repricer / dollar is Ki-magic which is **Free for new** sellers up to 50 SKUs.

45 Views

Figure F.9: Since Deleted Comment on Avoiding Price Wars on Quora.

Notes: This is a screenshot from a since-deleted comment on popular forum Quora.com that details several strategies for avoiding price wars on Amazon. The name of the commenter has been censored by the author.

G Additional Figures

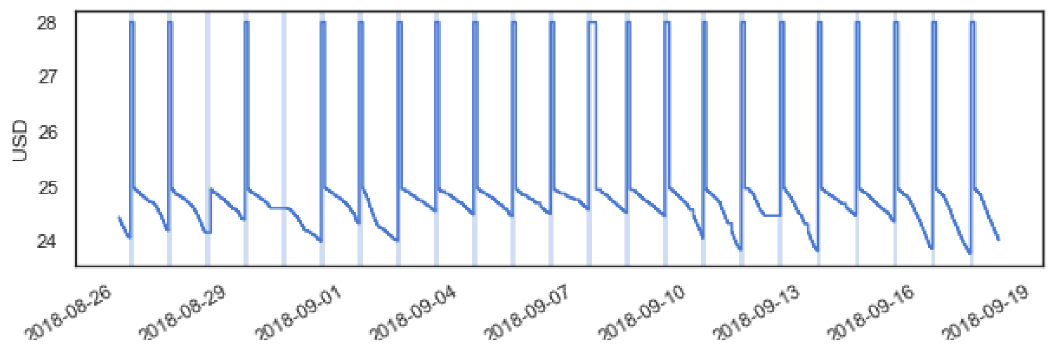

Figure G.1: Typical Day-Long Cycle

Notes: This figure shows the price (on the vertical axis) of a typical offer 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.

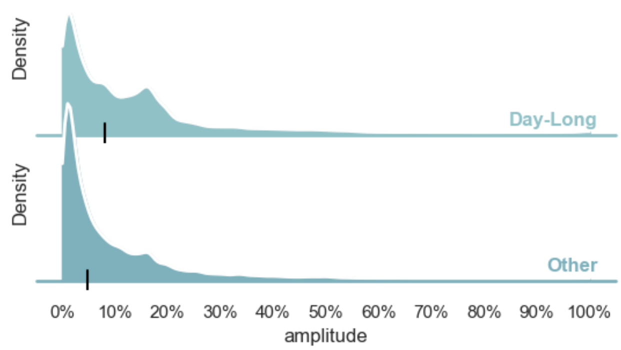

Figure G.2: Distribution of Cycling Amplitude

Notes: The figure shows the distribution (black bars: median) of cycle amplitude as a fraction of the maximum price separately for day-long cycles and other periods.