

Online Appendix

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

Leon Musolff

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*

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

These equilibria all result in identical price paths and payoffs.

Proof.

1. Suppose $\underline{p}_C \geq \bar{p}_C$. Then $p_C \equiv \bar{p}_C$ for all values of p_U . This cannot be part of an equilibrium:
 - (a) Suppose $\bar{p}_C > c$. Then, in equilibrium, U must have $\underline{p}_U \in [0, \bar{p}_C)$. Thus, C is making zero profits and would be better off by setting \bar{p}_C arbitrarily high (ensuring a cycle with positive profits).
 - (b) Suppose $\bar{p}_C = c$. Then, no matter what \underline{p}_U is played, C could deviate to playing $\bar{p}_C = 2\underline{p}_U$ 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 $\bar{p}_C > \underline{p}_C \geq \underline{p}_U$. 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 $(\underline{p}_C, \bar{p}_C) = (\ell^*, u^*)$. 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

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

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

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

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 $D > 0$. \square

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

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

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. \square

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

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

\square

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 $v_{\gamma N} = 0$. \square

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, $\alpha_U = 1$ and hence this inequality must be satisfied as $\epsilon \rightarrow 0$. 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 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*

$$\left((\bar{p}_1, \underline{p}_1), (\bar{p}_2, \underline{p}_2) \right) = \left((c, c), (c, c) \right).$$

Proof. We proceed by ruling out all other equilibria.

1. If $\underline{p}_i \geq \bar{p}_i > c$, then $p_i \equiv \bar{p}_i$ whence a profitable deviation is to set $\underline{p}_j = \bar{p}_j = \bar{p}_i - \epsilon$ for some small ϵ . Hence, this cannot be part of an equilibrium.
2. If $\underline{p}_i > \bar{p}_j > c$, then $p_j \equiv \bar{p}_j$ whence a profitable deviation is to set $\underline{p}_i = \bar{p}_i = \bar{p}_j - \epsilon$ for some small ϵ . Hence, this cannot be part of an equilibrium.
3. If $\bar{p}_i > \bar{p}_j > \underline{p}_i > \underline{p}_j$, then note that player i is choosing the lower end of the cycle and his choice \underline{p}_i must hence satisfy $\frac{\partial V_r(\underline{p}_i, \bar{p}_j)}{\partial \underline{p}_i} = 0$, i.e.

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

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

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

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

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

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

□

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 $k \rightarrow 0$, 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; the U agent (with unchanged strategy) will then (in the limit) match C's price.²¹

Let $D(p) = D_i(p, p) + D_j(p, p)$. In the limit as $k \rightarrow 0$, 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 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 off against (i)).²²

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 1b 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, it's 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 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)

²¹Recall that C can implement a fixed price p by setting $\bar{p} = p$ and $p > \bar{p}$.

²²Note 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 embarrassment 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, i.e. the period between a competitor’s price change and my own (next) price change never exceeds $10x$ the average since the start of the cycle,*
3. *the distance between maxima is at most 40% of the amplitude,*
4. *the distance between minima is unconstrained, and*
5. *there are at least 7 steps in the cycle.*

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

To illustrate which type of behavior this definition captures, we provide an example in Figure D.1. The blue line indicates the landed price for a given offer, and a green backdrop highlights




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.


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 1.49% of the time. However, conditional on at least one cycling offer, the probability that there are two cycling offers is 35.31%. 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. For a more detailed analysis of how often the algorithm detects cycling by the number of offers, see Figure D.2.

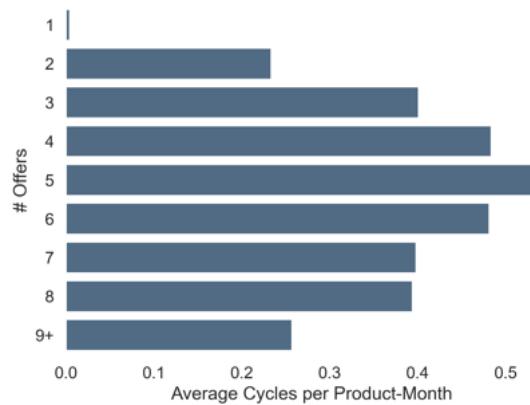
Price War Recognition. We utilize the following definition of a price war:

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

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



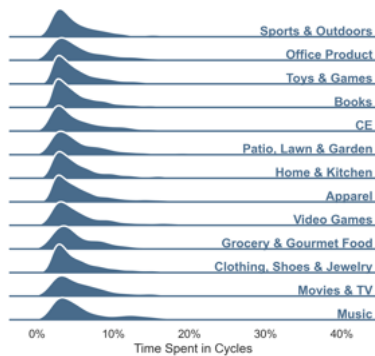
(a) Cycling Propensity by # Offers.



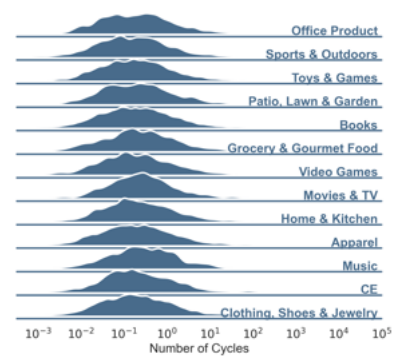
(b) Average # Cycles by # Offers.

Figure D.2: Cycling Prevalence by Number of Offers.

Notes: This figure illustrates the cycling propensity and average cycles per product-month by number of offers. The left panel shows the cycling propensity, which is the fraction of time that the offer is cycling. The right panel shows the average cycles per product-month, which is the average number of cycles per product-month.



(a) Distribution Of Fraction If Not Zero



(b) Distribution Of Count If Not Zero

Figure D.3: Distribution of Cycling Prevalence by Product Category.

Notes: This figure contains the distributions underlying the averages in Figure 11. The unit of observation is a product-month.

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.)

	(1)	(2)	(3)	(4)	(5)	(6)
	Log # Unique Prices	Log Price Span	Log Price	Log Mean Min.	Log Buybox Price	Buybox Share
Days Since Repricer Activ.						
-15 Days	0.033* (0.020)	0.002 (0.002)	-0.000 (0.008)	-0.010 (0.007)	-0.006 (0.007)	-0.005 (0.015)
-5 Days	0.029** (0.012)	0.002 (0.002)	0.001 (0.003)	-0.006* (0.003)	-0.005 (0.003)	0.009 (0.007)
0 Days	0.465*** (0.018)	0.026*** (0.002)	-0.071*** (0.005)	0.001 (0.003)	0.001 (0.003)	0.057*** (0.008)
5 Days	0.381*** (0.018)	0.016*** (0.002)	-0.092*** (0.005)	-0.015*** (0.004)	-0.014*** (0.004)	0.099*** (0.009)
15 Days	0.315*** (0.020)	0.012*** (0.002)	-0.127*** (0.006)	-0.046*** (0.005)	-0.046*** (0.005)	0.113*** (0.011)
30 Days	0.247*** (0.022)	0.009*** (0.002)	-0.147*** (0.007)	-0.065*** (0.006)	-0.062*** (0.006)	0.095*** (0.012)
60 Days	0.200*** (0.031)	0.010*** (0.003)	-0.167*** (0.010)	-0.097*** (0.009)	-0.095*** (0.009)	0.091*** (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 E.1: 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 6. 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 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.

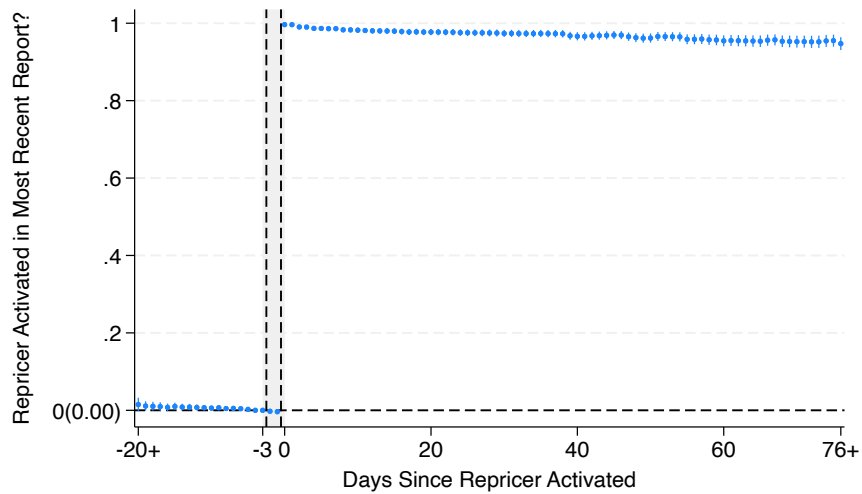


Figure E.1: Stickyness of Repricer Activation.

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 6 and Table E.1.

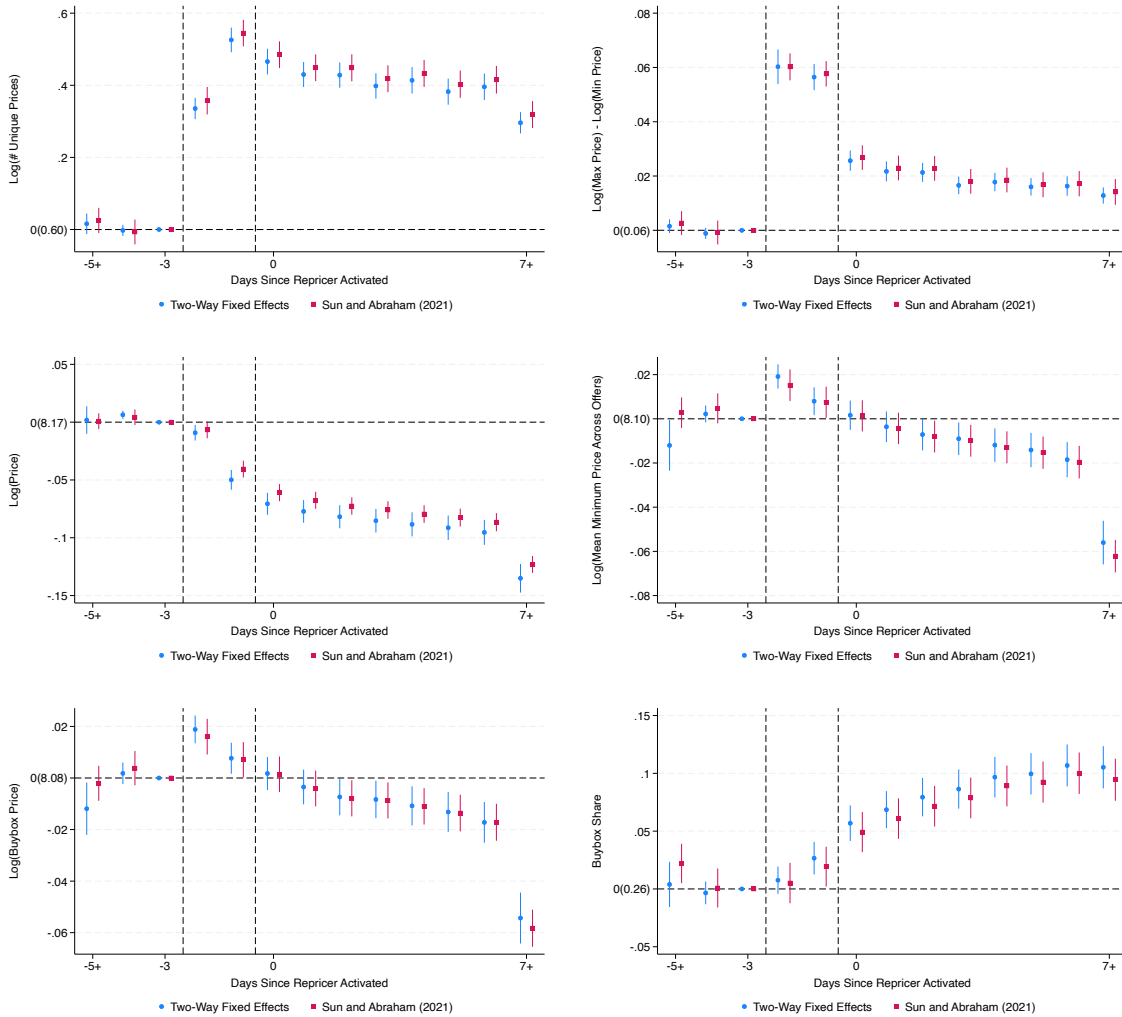


Figure E.2: Reprinter Activation Event Study Is Robust To Alternative Estimators.

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.

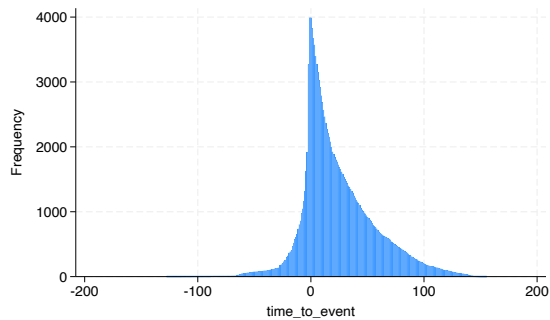


Figure E.4: Distribution of Time To Repricer Activation.

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.

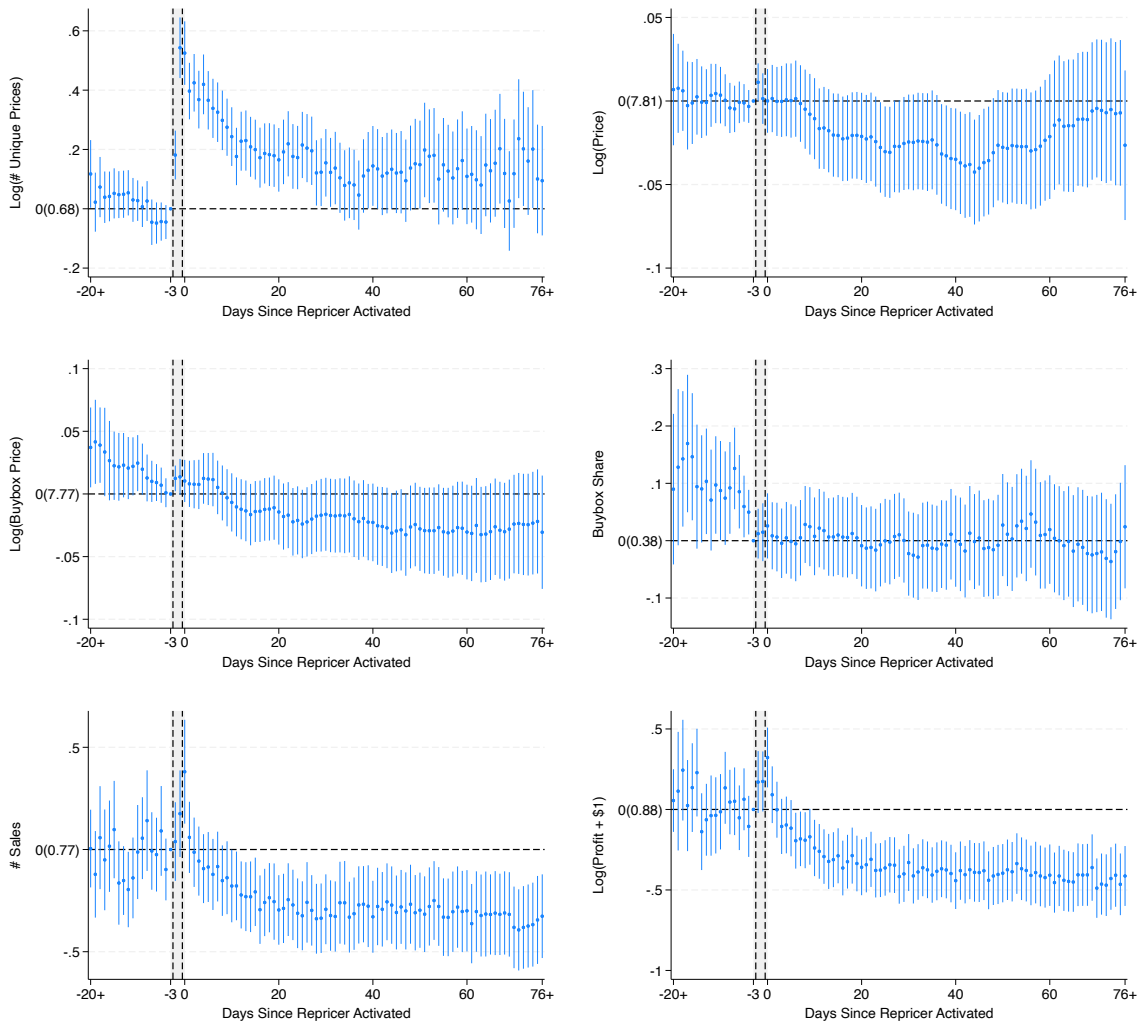


Figure E.5: Repricer Activation on Offers Reporting Sales And Cost Data.

Notes: This figure exhibits results from an event study that measures the effects of initial repricer activation; it differs from 6 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.

	(1)	(2)	(3)	(4)	(5)	(6)
	Own Log # Unique Prices	Own Log Max Price	Comp. Log Max Price	Own Log Min Price	Market Log Buyb. Price	Market Log Min Price
Days Since Resetting Activ.						
-90 Days	-0.166** (0.069)	0.001 (0.015)	-0.029 (0.024)	0.010 (0.016)	-0.012 (0.013)	-0.017 (0.012)
-60 Days	-0.054 (0.055)	0.001 (0.011)	0.003 (0.020)	-0.028** (0.014)	-0.012 (0.009)	-0.020** (0.008)
-30 Days	-0.014 (0.045)	0.010 (0.008)	-0.003 (0.006)	-0.002 (0.007)	-0.000 (0.006)	-0.003 (0.005)
0 Days	0.393*** (0.063)	0.125*** (0.017)	-0.011 (0.007)	0.011 (0.010)	-0.000 (0.007)	0.011 (0.010)
30 Days	0.441*** (0.072)	0.145*** (0.021)	0.004 (0.013)	0.041** (0.017)	0.031** (0.015)	0.028* (0.015)
60 Days	0.400*** (0.079)	0.167*** (0.021)	0.034** (0.014)	0.065*** (0.019)	0.045*** (0.016)	0.040** (0.016)
90 Days	0.293*** (0.092)	0.181*** (0.022)	0.078*** (0.017)	0.106*** (0.021)	0.081*** (0.020)	0.080*** (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 E.2: 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 8. 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 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.

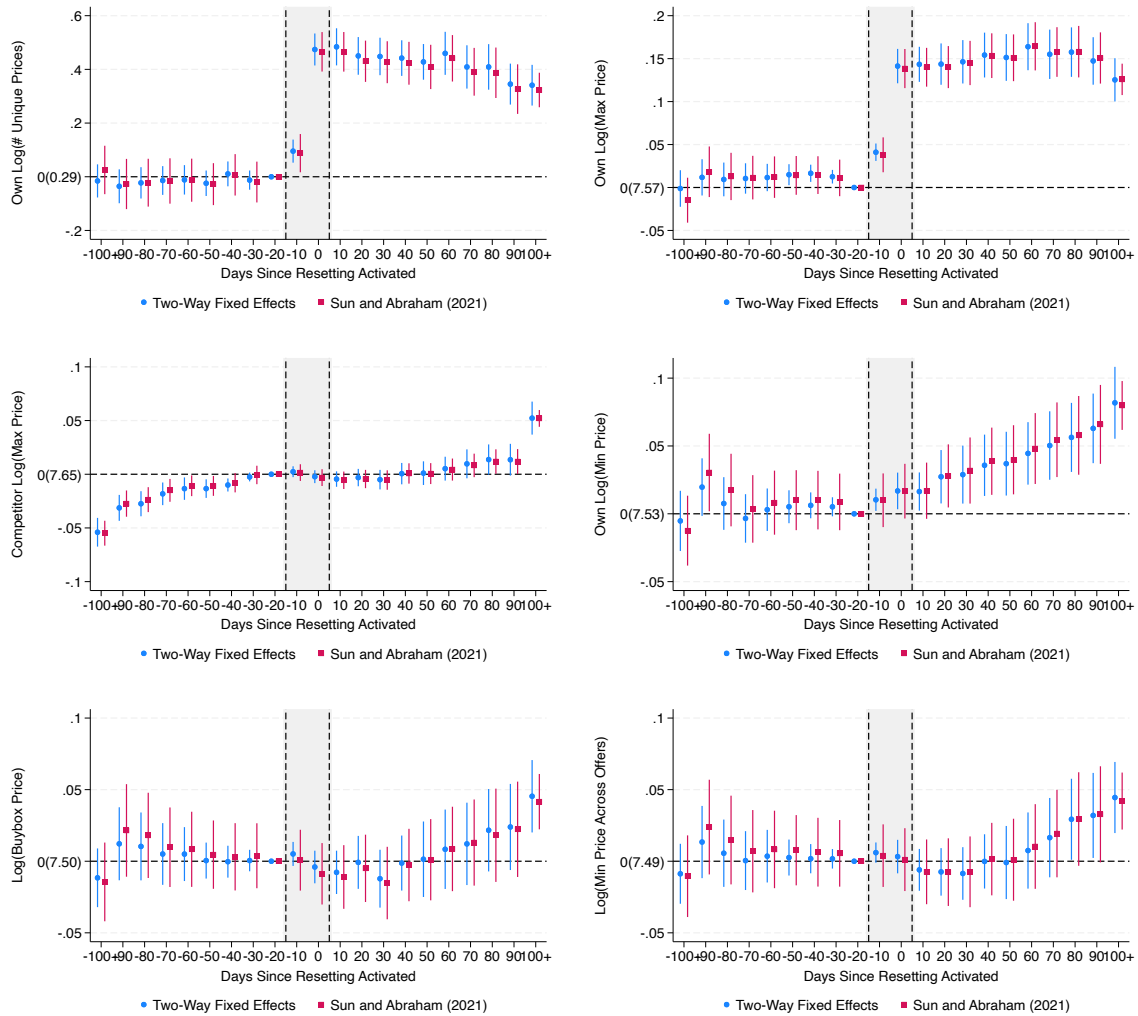


Figure E.6: Resetting Effect on Prices (Including Crowded Markets).

Notes: This figure exhibits results from an event study that measures the effects of activating a resetting strategy; it is identical to Figure 8 but includes all products, not just those between two and five competitive offers. The horizontal axis measures days since the resetting strategy was activated for the source merchant’s offer (with 0 being the day of activation). By contrast to Figure 8, we find a similar (if slightly smaller) effect of resetting activation on outcomes that indicate continued deployment of a resetting strategy (e.g., the own maximum price increases even more clearly and discontinuously at the event time). However, as would be expected, we find weaker and more ambiguous effects on market-level outcomes and on outcomes for competitors. This is consistent with the presence of many non-repricing competitors on these crowded markets. Nevertheless, the effects on the own minimum price are still positive and statistically significant.

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,

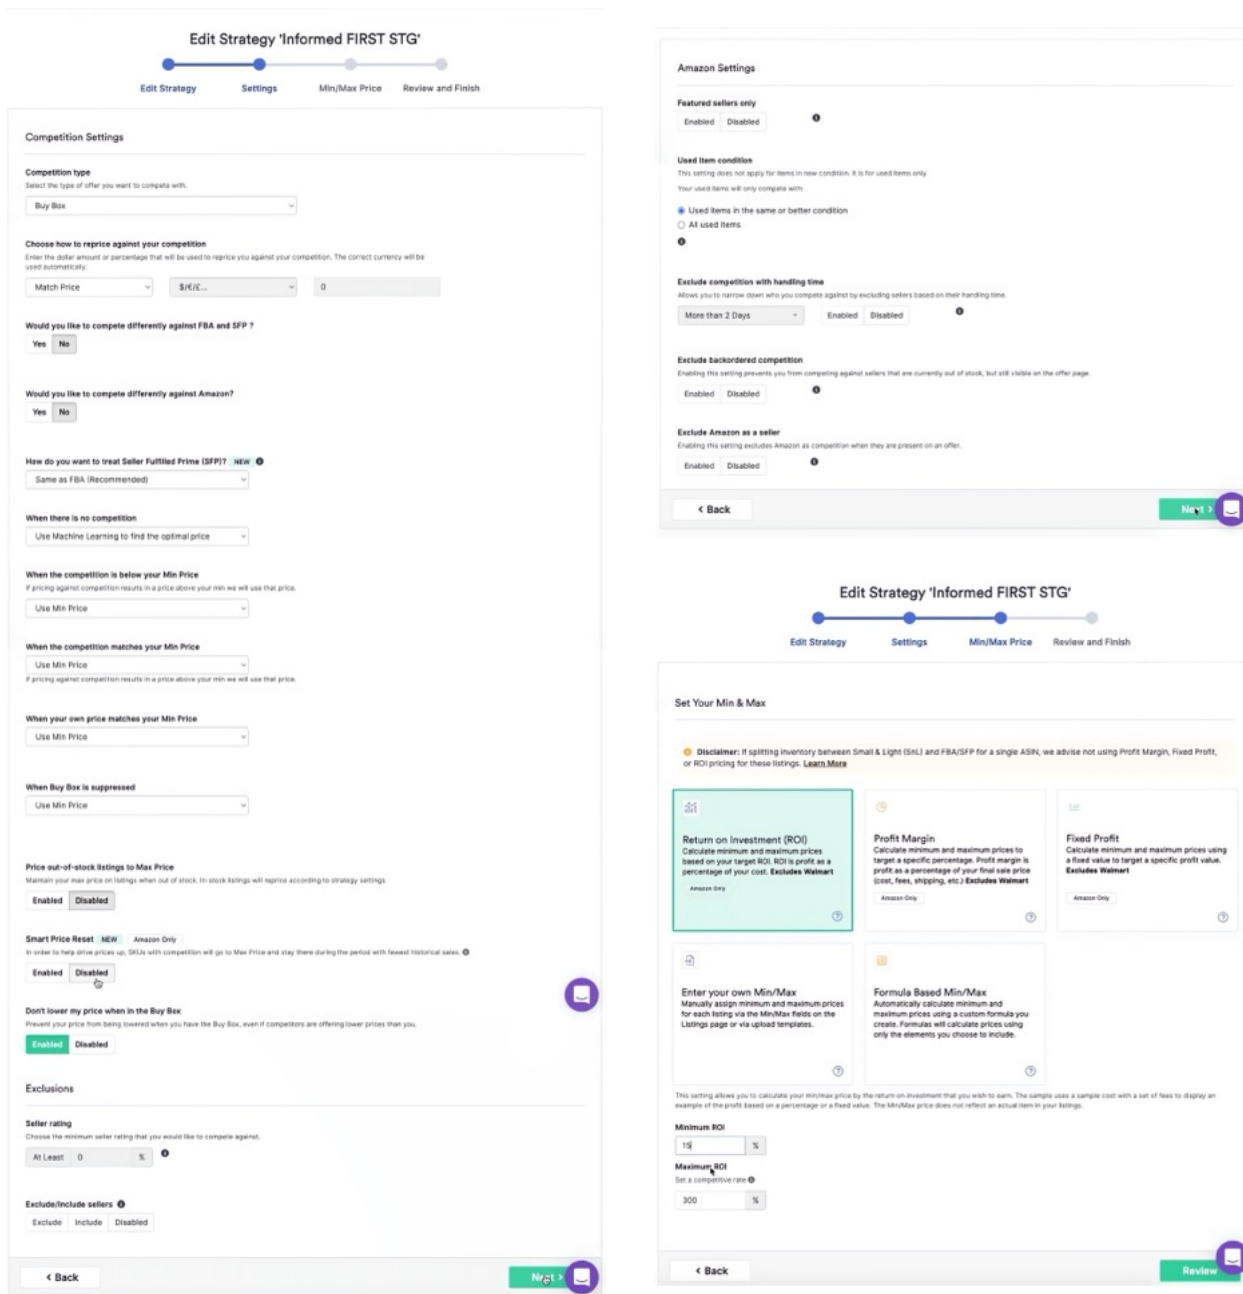


Figure F.1: Repricing Interface: Informed.

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 4) 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 4) 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.

AURA Strategy Builder Need Help? X

Marketplace Strategy Type **Competitors** Pursuit Maintenance Min/Max Pricing

Your Competitors

Define how Aura identifies your competitors.

What's the Competitors section?

Only compete with sellers who meet the following criteria.

Competition type
 Lowest Overall
 The type of competition that your listings will compete against.

Featured sellers only

 Featured sellers are those that are eligible for the buy box.

Exclude Amazon as a competitor

 Choose to ignore Amazon when found on a listing.

Exclude/Include sellers by their ID Coming Soon
 Include Exclude

Choose to exclude or include sellers regardless of other settings.

Exclude sellers by their feedback rating
 80 %
 Sellers that meet or exceed this feedback level will be considered competitors.

Exclude back-ordered sellers

 Exclude sellers whose inventory is not yet available.

Exclude competition with longer handling time

 More than 3 days
 Competitors with a long handling time will be excluded.

Your used items will compete against
 Used items in the same or better condition
 The subcondition of used items to be considered competition.

Marketplace Strategy Type Competitors **Pursuit** Maintenance Min/Max Pricing

Getting the Buy Box

Customize how Aura competes for the Buy Box.

What's the Pursuit section?

Adjust the actions that Aura will take to get you into the Buy Box.

How to price against the competition
 Price below by \$ 0.01
 The amount Aura will adjust your price by when competing with other sellers.

Price differently against Prime sellers (FBA and SFP)

Price differently against Amazon

 Match price
 The amount Aura will adjust your price by when competing with Amazon.

When the Buy Box is suppressed
 Customize price
 Set to Current price - \$ 0.02
 Action for Aura to take when no competition is found on a listing. In the case of a "Only Buy Box" competition type, this would indicate a suppressed Buy Box.

When the competition is below your min price
 Use min price
 Action for Aura to take when the competition is below your min price.

When the competition matches your min price
 Use min price
 Action for Aura to take when the competition is at your min price.

When your price reaches your min price
 Use min price
 Action for Aura to take when your price has reached your minimum.

Use max price when out of stock

 Listings can be set to max price when going out of stock.

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

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

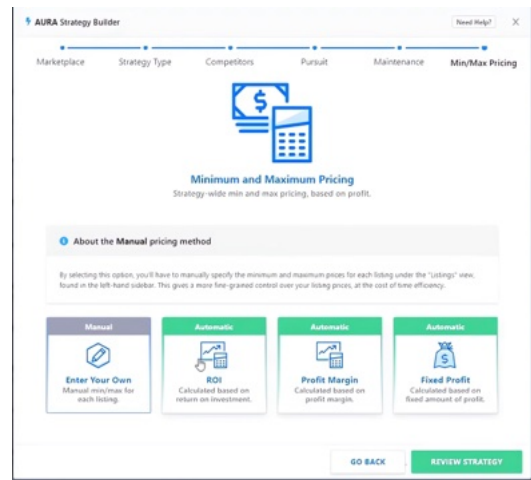
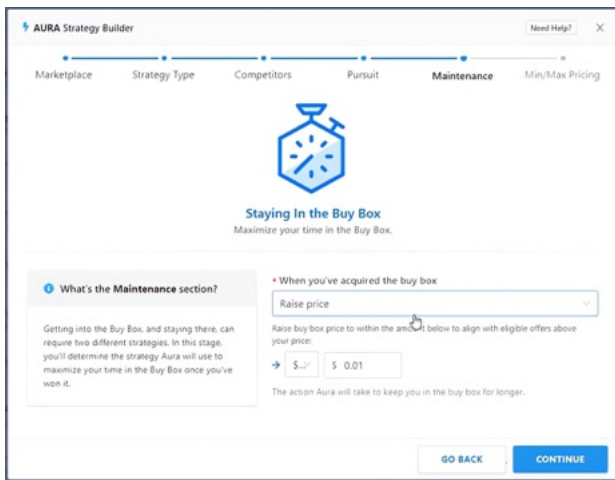


Figure F.3: Repricing Interface: Aura (Part II).

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

The figure consists of two side-by-side screenshots of the BQool repricing interface, showing the 'Custom Rule' configuration process.

Left Screenshot: Competitors Step

- Progress Bar:** Rule Type (selected), Competitors, Get Buy Box, Stay in Buy Box, Schedule, Rule Overview.
- Section Header:** Custom Rule: Define how BQool identifies your competitors.
- Competitors Section:**
 - Choose who you would like to compete with when you're not in the Buy Box. Compete with: Buy Box Price.
 - Choose Competitor(s):
 - Amazon
 - Non-Featured FBA
 - FBA
 - Non-Featured FBM
 - FBM
- Custom Settings Section:**
 - Filter the competitors by the settings below:
 - Set my SFP offers as: FBM
 - Set competitor's SFP offers as: FBM
 - Exclude Sellers by ID: off
 - Exclude Sellers with free shipping: off
 - Exclude Sellers with expedited shipping: off
 - Exclude Sellers Positive Feedback Rating lower than: on, 65.00 %
 - Exclude Sellers Total Feedback Count lower than: on, 0
 - Exclude Sellers with Backordered product: on
 - Exclude Sellers with handling time more than: off, 3 days
 - Set the item condition you would like to compete against: off, Same
- Buttons:** Back, Cancel, Save and Finish, Next.

Right Screenshot: Get Buy Box Step

- Progress Bar:** Rule Type, Competitors, Get Buy Box (selected), Stay in Buy Box, Schedule, Rule Overview.
- Section Header:** Custom Rule: Define how BQool reprices when you're not in the Buy Box.
- Get Buy Box Section:**
 - When you're not in the Buy Box or "Stay in Buy Box" setting has been disabled, these are the settings that will be used.
 - Buy Box Winner found between Min & Max:**
 - When Buy Box Winner is above Min Price: Buy Box Price, 0.00 \$
 - When Buy Box Winner is below Min Price: Use Auto Compare (1 Double Buy Box Only)
 - When Buy Box Winner equals Min Price: Use Min Price
 - Buy Box Winners found outside Min & Max:**
 - When Buy Box Winners found above Max Price: Max Price, 0.00 \$
 - When Buy Box Winners found below Min Price: Min Price, 0.00 \$
 - When Buy Box Winners found below Min Price & above Max Price: Min Price, 0.00 \$
 - Buy Box Winner not found:**
 - When no one owns the Buy Box: Min Price, 0.00 \$
 - Competition not found:**
 - When you are the only seller: Use Max Price
 - When Buy Box Winners excluded by your settings: Do Not Reprice
 - Price Change Safety Net:**
 - Set the maximum amount your price can decrease for one price change:
 - off: When adjusted price is lower than Your Price by: 0.00 %
 - Then limit the adjusted price to = Your Price: 0.00 %
 - Min Price Protection Settings:**
 - When Adjusted Price equals or is below Min Price: Use Min Price
 - Advanced Settings:**
 - Define further how you want to compete against specific competitors. Note: When [Advanced Settings] are turned on, [Get Buy Box] will be overridden.
 - FBA vs Amazon: off
 - FBA vs FBA: on
 - FBA vs FBM: off
 - FBM vs Amazon: off
 - FBM vs FBA: on
 - FBM vs FBM: on
- Buttons:** Back, Cancel, Save and Finish, Next.

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

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

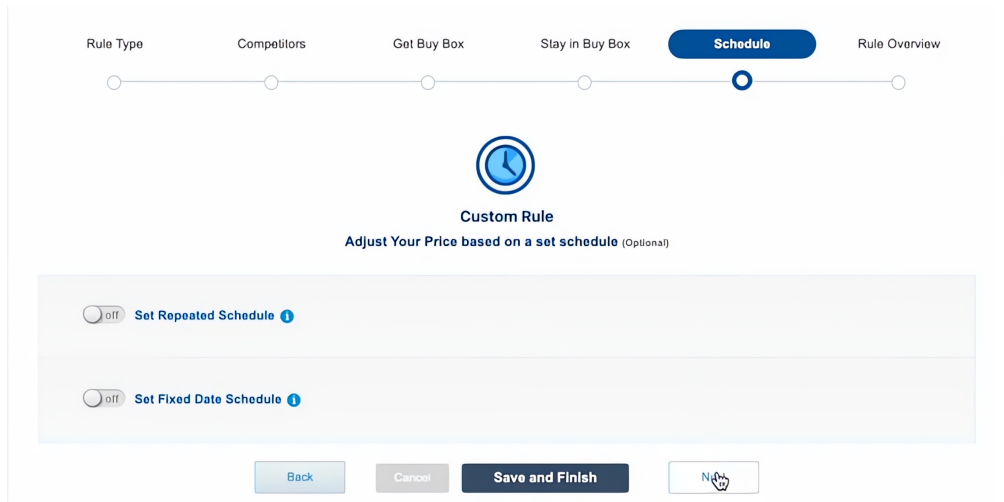
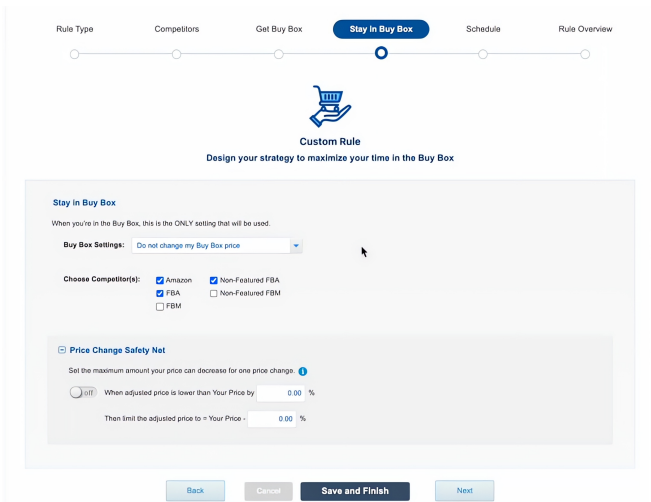


Figure F.5: Repricing Interface: BQool (Part II).

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

Repricing Rules > Add New Repricer

Basic Info Advanced Options Create Repricer

Repricer Name * Repricer 2

Is Active?

Default Pricing Rule Below the competition by Amount 0

Manually assign min/max

MIN Price No Minimum above

MAX Price Net Margin Value above

Repricing Rules > Add New Repricer

Basic Info Advanced Options Create Repricer

If BuyBox winner? * The repricer will carry on as normal, repricing up and down

Used listings * Used products compete equally against all other used listings

Featured merchants * Compete against all merchants

Out of bounds sellers * Compete against all merchants

Ignore the following merchants

If no competitors * Change my price to my highest allowed price

Ignore rule? * No, always apply the default rule if none of my competitor rules have been matche

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

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

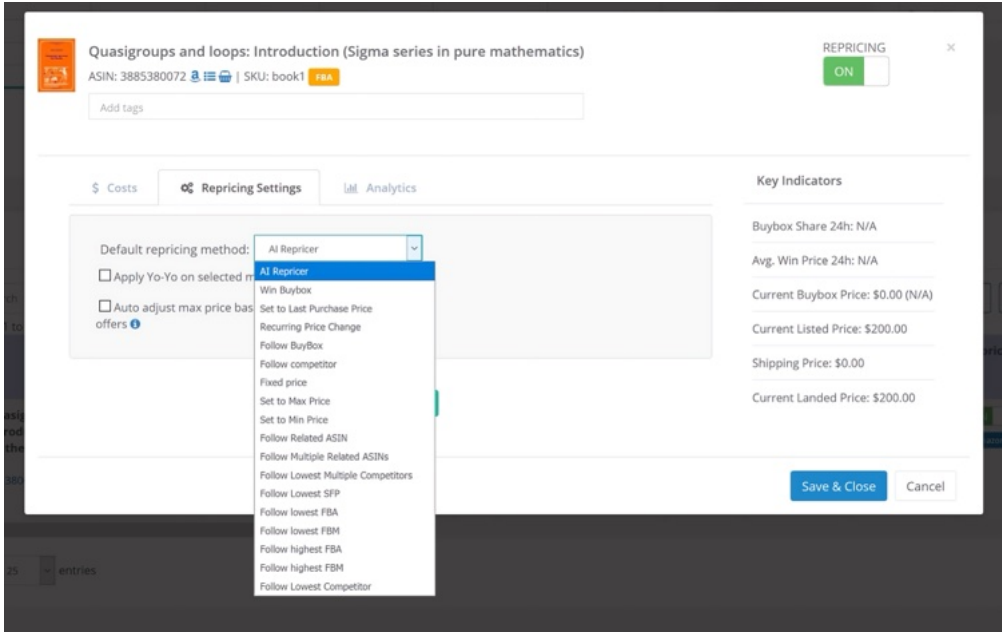


Figure F.7: Repricing Interface: SellerSnap.

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

Yo-Yo Repricing Rule



Yo-Yo repricing allows you to manually set a loop increasing the price to Max and then reverting to an automatic repricing method such as Game Theory Repricing.

Yo-Yo can be added to any repricing method by selecting "Apply Yo-Yo on Selected Method":

Default repricing method: AI Repricer Apply Yo-Yo on selected method

Increase price to max price every 180 minutes. Stay at max price for 20 minutes.

Profit Margin Optimized ⓘ

Add Condition

In this example, the repricer will increase the price to Max every 180 minutes, stay at Max for 20 minutes and then revert to AI repricing.

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 amazon's 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.

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

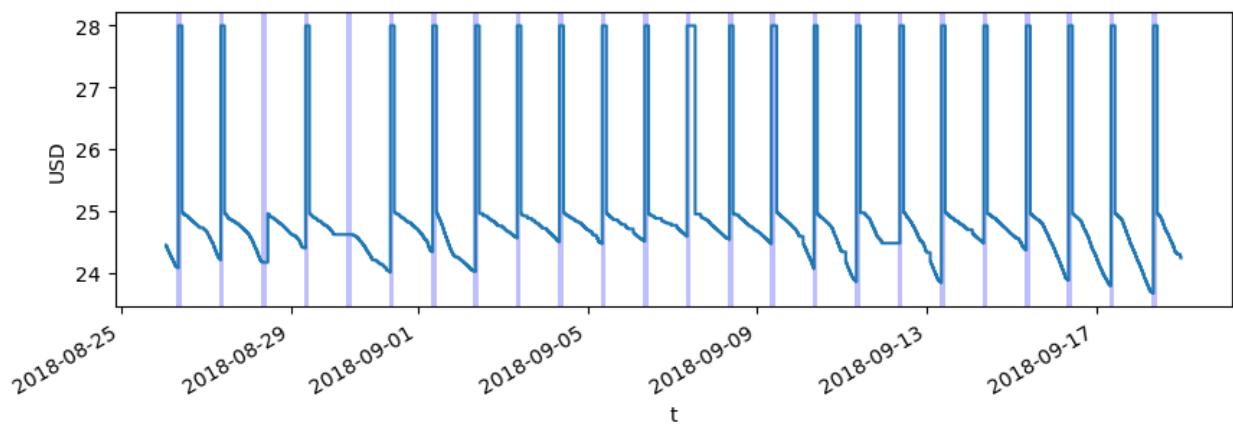


Figure G.1: Typical Day-Long Cycle

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.

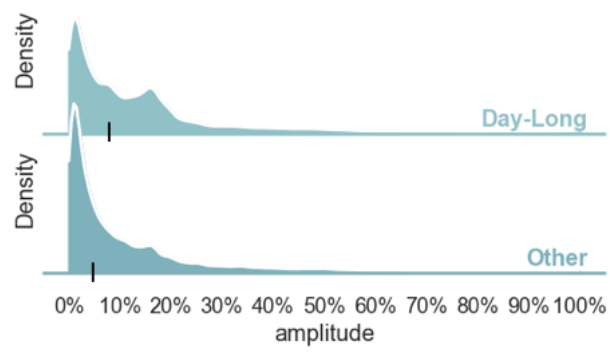


Figure G.2: Distribution of Cycling Amplitude

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.