Online Appendix for "Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence"

Christopher T. Conlon*
Julie Holland Mortimer†

June 14, 2021

A. Theoretical Motivation for Offering an AUD

A.1. Foreclosure and Optimal Assortments: A Motivating Example

We define the difference in payoffs between two assortments as $\Delta\pi(a, a') = \pi(a) - \pi(a')$. We introduce the possibility that the dominant firm M offers the retailer a lump sum transfer T in exchange for switching from assortment a' to assortment a. For this to be an equilibrium the following necessary conditions must be met:

$$\Delta\pi^R + T \geq 0 \quad \text{(Retailer IR)}$$
$$\Delta\pi^M - T \geq 0 \quad \text{(Mars IR)}$$

The retailer must prefer to receive the rebate under assortment a than to not receive the rebate under assortment a'. Meanwhile the dominant firm must prefer to pay the rebate under assortment a over not paying the rebate under assortment a'. For a to represent an equilibrium assortment, it must also be the case that no player has an incentive to deviate, including the rival firm H. Were H to offer its own transfer T_h in exchange for the retailer choosing assortment a' instead of a this becomes the opposite of the Mars IR constraint:

$$\Delta\pi^H + T_h \leq 0 \quad \text{(Hershey Deviation)}$$

*Department of Economics, NYU Stern, 44 W 4th St, New York City, NY 10012, and NBER. email: cconlon@stern.nyu.edu

†Department of Economics, Boston College, 140 Commonwealth Ave., Chestnut Hill, MA 02467, and NBER. email: julie.mortimer.2@bc.edu

We can consider the 'bidding for representation' argument of Bernheim and Whinston (1998), where each transfer is set at the maximum amount so that $T_h = -\Delta\pi^H$ and $T = \Delta\pi^M$ in order to see whose transfer persuades the retailer:

$$\begin{aligned}\pi^R(a) + T &\geq \pi^R(a') + T_h \\ \Delta\pi^R + \Delta\pi^M &\geq -\Delta\pi^H \\ \Delta\pi^R + \Delta\pi^M + \Delta\pi^H &\geq 0 \qquad \text{(Three-Party Surplus)}\end{aligned}$$

This tells us if the three conditions are satisfied (Retailer IR, Mars IR, and Three-Party Surplus) then some transfer T (conditioned on assortment a) makes a an equilibrium when the no-transfer equilibrium is a'. In the subsequent section we show how the AUD contract allows the dominant firm to design the rebate threshold $\overline{\pi}^M$ to pay the transfer conditional on particular assortments a.[1]

We show how to adapt this setup to our empirical example. There are three potential assortments for the last two products on the shelf, two Mars products (M, M), two Hershey's products (H, H), or the best of each (H, M). Each manufacturer earns higher profits when more of their own products are stocked. Absent transfers, the retailer prefers to stock more Hershey's products and fewer Mars products. We assume that the profits of each agent can be ordered as follows (this mimics the actual payoffs in our empirical example):

$$\begin{aligned}\pi^R(H,H) >& \pi^R(H,M) > \pi^R(M,M) \\ \pi^H(H,H) >& \pi^H(H,M) > \pi^H(M,M) \\ \pi^M(M,M) >& \pi^M(H,M) > \pi^M(H,H) \qquad (1)\end{aligned}$$

Given the ordering of profits above, absent the rebate the retailer prefers the assortment (H, H). Now we can consider decomposing profit differences into two steps. The first is the difference between (H, H) and (H, M) which we call Δ_H and the second is the difference between (H, M) and (M, M) which we call Δ_M so that $\Delta = \Delta_H + \Delta_M$ represents the difference between (H, H) and (H, M).

Conditions A $a = (M, M)$ and $a' = (H, H)$. $\Delta\pi^R + T \geq 0$ (IRR), $\Delta\pi^M - T \geq 0$ (IRM) and $\Delta\pi^R + \Delta\pi^M + \Delta\pi^H \geq 0$ (3 Party).

Conditions B $a = (M, H)$ and $a' = (H, H)$. $\Delta_H\pi^R + T \geq 0$ (IRR), $\Delta_H\pi^M - T \geq 0$ (IRM)

[1] We also show how it can be used to select effort levels e in accordance with the (IC) constraint of the retailer.

and $\Delta_H \pi^R + \Delta_H \pi^M + \Delta_H \pi^H \geq 0$ (3 Party).

Conditions C $a = (M, M)$ and $a' = (M, H)$. $\Delta_M \pi^R + T \geq 0$ (IRR), $\Delta_M \pi^M - T \geq 0$ (IRM) but not necessarily the three-party surplus condition.

If conditions A hold then we have shown that there exists a transfer T such that (M, M) is an equilibrium as no player possesses a profitable deviation. It is also the case that the three-party surplus or industry profits $\pi^I = \pi^M + \pi^H + \pi^R$ are higher under (M, M) than (H, H) as $\Delta \pi^I \geq 0$.

From conditions B we know that that $\Delta_H \pi^I \geq 0$ or that the three-party surplus under (H, M) is higher than that under (H, H).

It could be that $\Delta_M \pi^I < 0$ or that the (H, M) assortment rather than the (M, M) assortment maximizes the three-party surplus. This does not contradict any of the other conditions.

The main takeaway is that M can set the transfer payments in order to obtain full (M, M) or partial (H, M) foreclosure. We show that under (A), full foreclosure is feasible. However, if (B), (C), and $\Delta_M \pi^I < 0$ also hold, full foreclosure does not lead to the assortment that maximizes overall industry surplus. In this case, partial foreclosure maximizes industry surplus, but full foreclosure leads to higher bilateral surplus among the retailer and dominant firm. As long as the dominant firm chooses the transfers and conditions, full foreclosure will be the equilibrium outcome.

The intuition behind this result relates to that of the *Chicago Critique* of Bork (1978) and Posner (1976), which we interpret as asking "When foreclosure is obtained in equilibrium, must the assortment necessarily be optimal?" Our answer is related to the work by Whinston (1990) on tying. When the dominant firm is able to condition the transfer payment on the (M, M) outcome, he can commit to tying the products together, and thus the equilibrium assortment need not maximize the surplus of the entire industry.

A.2. Effort Derivation

Consider the effort choice of the retailer faced with an AUD contract from (??):

$$\max_{(a,e)} \pi(a,e) = \begin{cases} \pi^R(a,e) + \lambda \cdot \pi^M(a,e) & \text{if } \pi^M(a,e) \geq \overline{\pi}^M \\ \pi^R(a,e) & \text{if } \pi^M(a,e) < \overline{\pi}^M. \end{cases}$$

It is helpful to temporarily ignore the assortment choice a and focus on effort only. In the case where the rebate is paid, we can express the retailer's problem as:

$$e_1 = \arg\max_e \pi^R(e) + \lambda \pi^M(e) \quad \text{s.t.} \quad \pi^M(e) \geq \overline{\pi}^M$$

The solution to the constrained problem is given by:

$$e_1 = \max\{e^R, \overline{e}\} \quad \text{where} \quad \overline{e} \text{ solves} \quad \pi^M(\overline{e}) = \overline{\pi}^M$$

If the rebate is not paid then:

$$e_0 = e^{NR} = \arg\max_e \pi^R(e)$$

The retailer's IC constraint:

$$\pi^R(e_1) + \lambda \pi^M(e_1) \geq \pi^R(e_0) \tag{IC}$$

and the dominant firm M's IR constraint:

$$(1-\lambda)\pi^M(e_1) \geq \pi^M(e_0) \tag{IRM}$$

When we consider the sum of (IC) and (IRM) it is clear that a rebate which induces effort level e_1 must increase bilateral surplus relative to e_0:

$$\pi^R(e_1) + \pi^M(e_1) \geq \pi^R(e_0) + \pi^M(e_0)$$

This provides an upper bound on the effort that can be induced by the rebate contract.

Thus, for $\overline{e} \geq e^R$, M can set the effort level of the retailer via the threshold $\overline{\pi}^M$, subject to satisfying the retailer's IR constraint. That is, the retailer must prefer to collect the rebate to their next best no rebate alternative (generally the (H, H) assortment).

A.3. Alternative Contracts

This section compares the AUD contract to other contractual forms; it is meant to be expositional and does not present new theoretical results.

Quantity Discount

A discount τ, can be mapped into λ (a share of M's variable profit margin). However the discount no longer applies to all q_m, only those units in excess of the threshold, so that $\rho(\overline{\pi}^M) = \max\left\{0, \frac{\pi^M - \overline{\pi}^M}{\pi^M}\right\}$. This implies $T \equiv \rho(\overline{\pi}^M) \cdot \lambda \cdot \pi^M$, so that as the threshold increases, M is limited in how much surplus it can transfer to R, assuming that the post-discount wholesale price is non-negative. In the limiting case, the threshold binds exactly and M cannot offer R any surplus. This makes the discount, rather than the threshold, the primary tool for incentivizing effort. (Recall that for the AUD, $\overline{e} \geq e^R$ implies that M can directly set the retailer's effort). This means that high effort levels, $e > e^R$, will be more expensive to the dominant firm under the quantity discount than under the AUD. In fact, the vertically-integrated level of effort is only achievable through the 'sell out' discount, where $\tau = w_m - c_m$ such that M earns no profit on the marginal unit, and some $\overline{q}_m$ significantly less than the vertically-integrated quantity.

Quantity Forcing Contract

The quantity forcing (QF) contract is similar to a special case of the AUD contract. Specify a conventional AUD $(w_m, \tau, \overline{q}_m)$ as:

$$\begin{cases} (p_m - w_m + \tau) \cdot q_m & \text{if } q_m \geq \overline{q}_m \\ (p_m - w_m) \cdot q_m & \text{if } q_m < \overline{q}_m \end{cases}$$

One can increase the wholesale price w_m by one unit, and the generosity of the rebate (τ) by one unit. Continuing with this procedure, the retailer profits when the threshold is met. For any $q_m \geq \overline{q}_m$, the retailer's profit remains unchanged, while its profit for any $q_m < \overline{q}_m$, tends to zero as $w_m \to p_m$. This has the effect of 'forcing' the retailer to accept a quantity at least as large as $\overline{q}_m$. By choosing the threshold, the QF contract can achieve the vertically-integrated level of effort, just like the AUD. For quantities $q_m > \overline{q}_m$, the AUD works like a QF contract plus a uniform wholesale price on 'extra' units.[2]. Without some

[2]For a more complete discussion of the connection between the AUD and the QF contract in the presence of a capacity constrained rival see Chao et al. (2018)

outside constraint on τ or w_m, and absent uncertainty about demand, the dominant firm has an incentive to increase τ and w_m together to replicate the QF contract.

Two-Part Tariff

One can also construct a two-part tariff (2PT), described by two terms: a share of M's revenue λ and a fixed transfer T from $R \rightarrow M$. The retailer chooses between the 2PT contract and the standard wholesale price contract.

$$\begin{cases} \pi^R(a,e) + \lambda \cdot \pi^M(a,e) - T & \text{if } 2PT \\ \pi^R(a,e) & \text{o.w.} \end{cases}$$

We define $\underline{\pi}^R = \max_{a,e} \pi^R(a,e)$ (the retailer's optimum under the standard wholesale price contract). For the retailer to choose the 2PT contract it must be that $\max_{a,e}\{\pi^R(a,e) + \lambda \cdot \pi^M(a,e) - T\} \geq \underline{\pi}^R$. An important case of the 2PT contract is the so-called 'sellout' contract where $\lambda = 1$. In this case, the retailer maximizes the joint surplus of $\pi^R + \pi^M$ and achieves the vertically-integrated assortment and stocking level. Just like in the AUD, this may lead to foreclosure of the rival H, even when that foreclosure is not optimal from an industry perspective. The dominant firm can choose T so that $\max_{a,e}\{\pi^R(a,e) + \pi^M(a,e)\} - T = \underline{\pi}_M$ and 'fully extract' the surplus from R. Likewise, the dominant firm can choose $T = (1 - \lambda_{AUD}) \cdot \overline{\pi}^M$ (the dominant firm's profits under the AUD) so long as the retailer is willing to choose the 2PT contract.

This indicates that it is also possible for a 2PT contract to implement the assortment and effort level that maximizes the bilateral profit between $M + R$, even if that assortment does not maximize overall industry profits. An important question is: how do the AUD and the 2PT differ? One possibility is that the AUD can be used to implement an effort level in excess of the vertically-integrated optimal effort, e^{VI}, which results in higher profits for M at the expense of the retailer. A major challenge of devising a 2PT in practice is arriving at the fixed fee T, especially when there are multiple retail firms of different sizes, and the 2PT contract (or menu of contracts) is required to be non-discriminatory.[3] It may be easier in practice to tailor sales thresholds to the size of individual retailers (as opposed to setting individual fixed-fee transfer payments).[4]

[3]Kolay et al. (2004) shows that a menu of AUD contracts may be a more effective tool in price discriminating across retailers than a menu of 2PTs. In the absence of uncertainty, an individually-tailored 2PT enables full extraction by M, but is a likely violation of the Robinson-Patman Act.

[4]Another possibility as shown by O'Brien (2013) is that the AUD contract can enhance efficiency under the double moral-hazard problem (when the upstream firm also needs to provide costly effort such as advertising).

B. Econometric Appendix

B.1. Additional Descriptive Figures

We provide alternative version of the descriptive figures in the text to illustrate how our experimental sample of machines that we use to estimate demand is similar (and different) from the overall enterprise of MarkVend. In all of these figures the unit observation is a machine-visit, and we average across machine-visits both by month. Thus machines that are visited more frequently are given more weight. We've also computed each of these figures re-weighted based on monthly machine sales and we obtain nearly identical results.

In Figure A1, we show the overall number of product facings in confections is relatively stable over time, but differs between our experimental sample of machines in white-collar office locations and MarkVend's wider enterprise which also includes some larger machines with more product facings in schools, museums, parks, etc. We should also note that if one considered an 'unbalanced' panel of MarkVend's entire enterprise, the number of product facings would appear to decline over time as the relative share of smaller office located machines grew relative to the share of larger machines.

In Figure A2, we reproduce the upper panel of Figure 2 from the main text for both the 'balanced' panel of 364 machines and our smaller experimental sample in office locations. In both cases, there is a pronounced shift around the beginning of 2008 when we believe that Mars changes the rebate threshold. Around the time the threshold changes, MarkVend replaces Chocolate Mars products with Chocolate Hershey products. This change takes place in both samples.

To further illuminate which product facings change over time we then produce Figure A3. Here we show that there are a set of base Mars products which are highly available in both samples, and don't vary much over time (Snickers, Peanut M&M's, Twix). We aggregate the two non-chocolate confections (Skittles and Starburst) as MarkVend tends to alternate machines (each machine stocks either Skittles or Starburst). A small number of machines (mostly in schools) stock both Skittles and Starburst, which explains why more than once facing is reported for the combined product.

Finally we reproduce the lower panel of Figure 2 from the main text in Figure A4. Again, the top pane is for a 'balanced' panel of 364 machines, while the bottom panel is for the 66 machines we use to conduct our experiments and estimate our demand model. Here we show the set of 'non base' product for each manufacturer. These are the products we generally view as competing for the final slots in the vending machine. The main takeaway is that when Mars reduces the threshold in 2008, MarkVend substitutes the worst performing Mars chocolate product (3 Musketeers) for the best performing Hershey product (Reese's Peanut Butter Cup). There are some other important differences between the two samples, the broader sample tends to include M&M Plain in roughly 80% of machines while our experimental sample includes it only 40%. The broader sample includes Sour Patch Kids (Hershey) in around half of machines, though they are almost never available in our office sample. Meanwhile the experimental sample stocks Raisinets (Nestle) in around 80% of machines as compared with 50% of the broader sample. We think these can largely be explained by differences in demand patterns between white collar office workers (Raisinets) and school-aged children (Sour Patch Kids), as well as the larger overall machines in the schools and museums.

Figure A1: Product Facings by Category

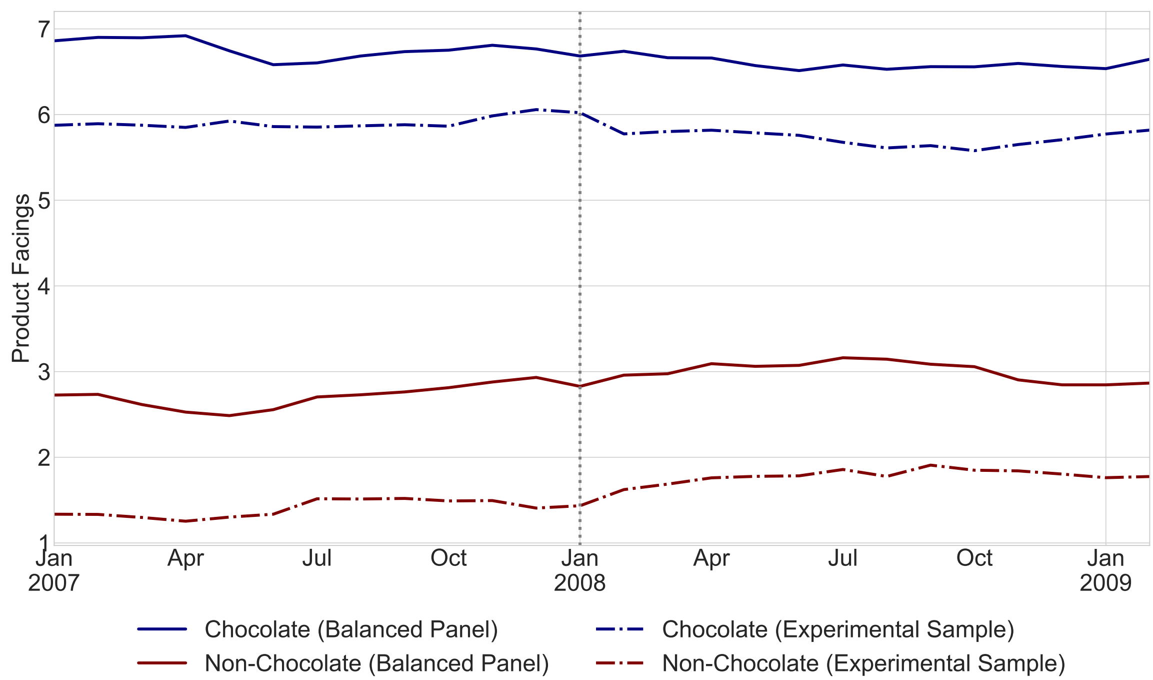

Notes: An observation is a machine-visit pair. Figure reports average product facings of confection products across machine-visits by month and product category for two sets of machines: a balanced panel of 364 MarkVend vending machines, as the set of 66 vending machines used in for our experimental product removals. Blue lines report chocolate confection products; red lines report non-chocolate confection products.

Figure A2: Product Facings by Manufacturer and Category

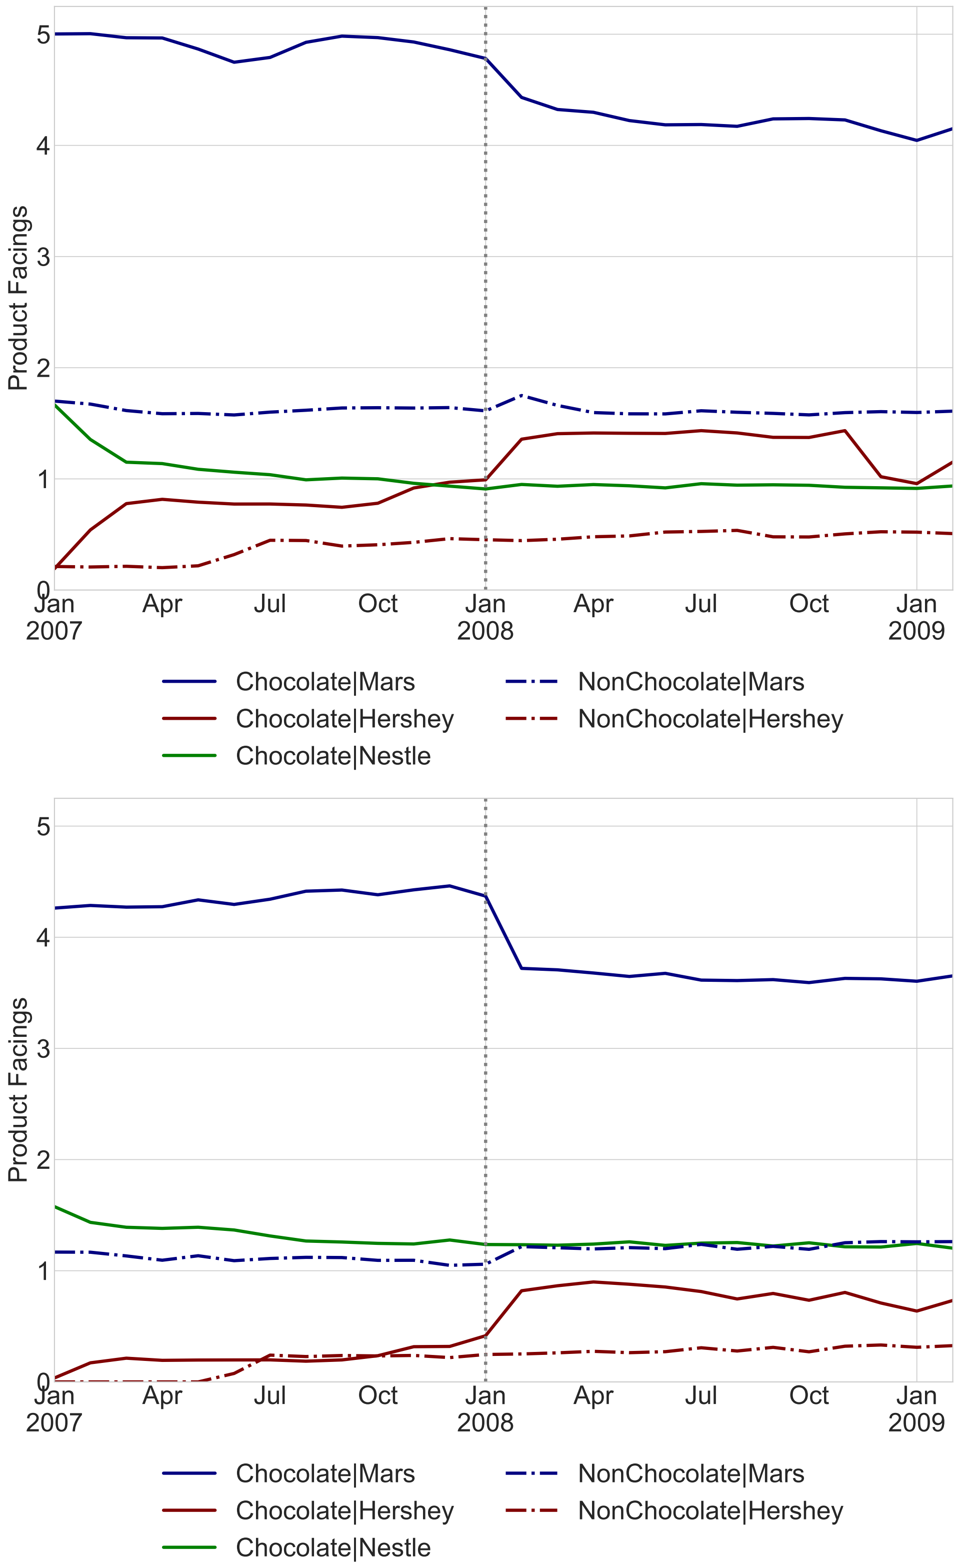

Notes: An observation is a machine-visit pair. Figure reports average product facings of confection products across machine-visits by month, product category, and manufacturer for two sets of machines: a balanced panel of 364 MarkVend machines (top pane), and the set of 66 machines used in for our experimental product removals (bottom pane).

Figure A3: Product Facings for Commonly Stocked (Base) Assortment

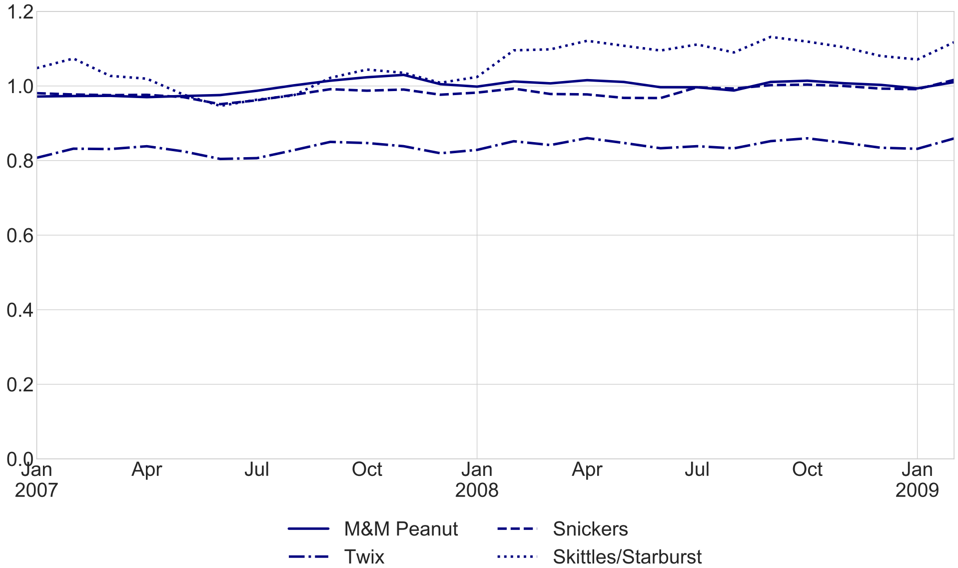

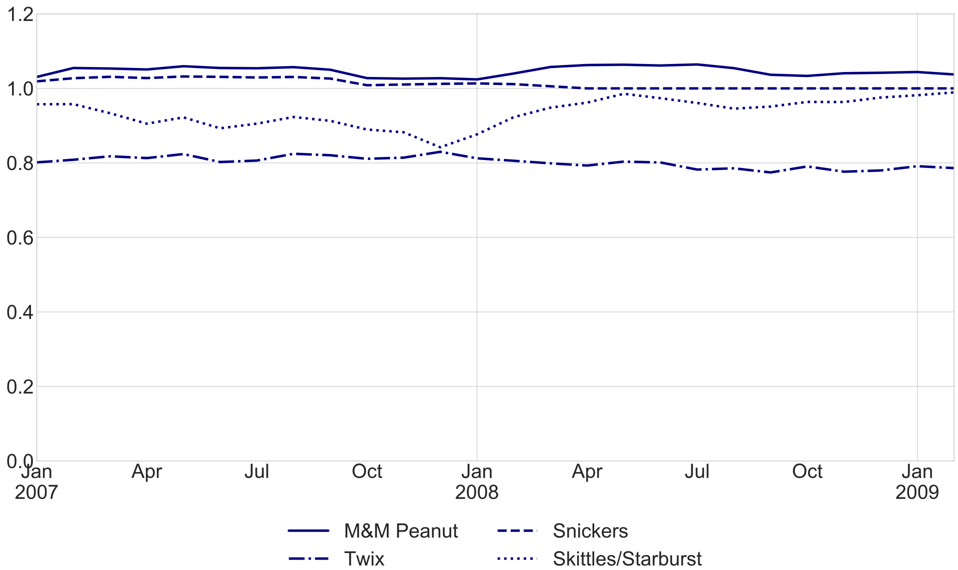

Notes: An observation is a machine-visit pair. Figure reports average product facings of products commonly included in MarkVend's base assortment across machine-visits by month for two sets of machines: a balanced panel of 364 MarkVend vending machines, as the set of 66 vending machines used in for our experimental product removals.

Figure A4: Product Facings for Marginal Products

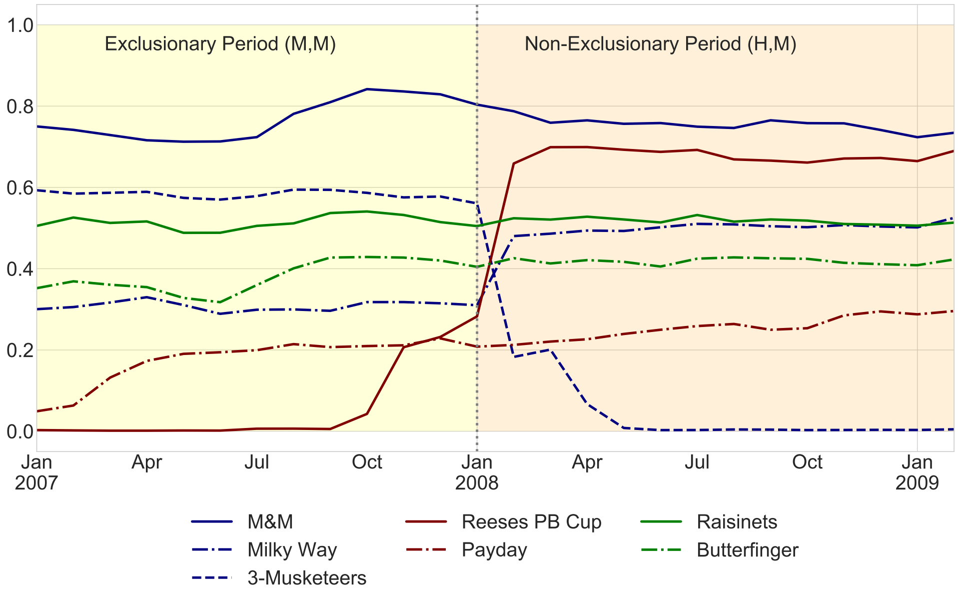

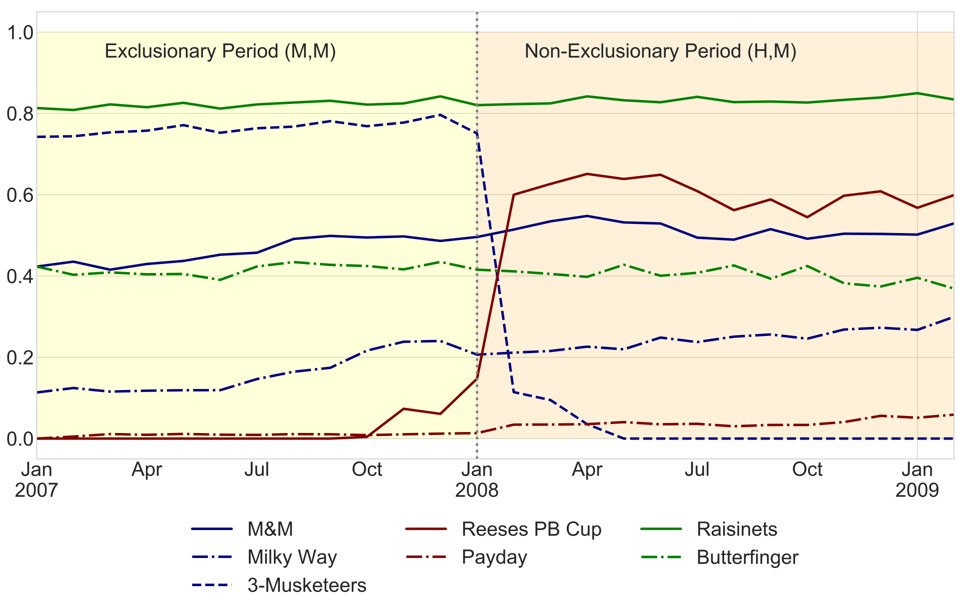

Notes: An observation is a machine-visit pair. Figure reports average product facings of *marginal* products in MarkVend's base assortment across machine-visits by month for two sets of machines: a balanced panel of 364 MarkVend vending machines, as the set of 66 vending machines used in for our experimental product removals.

B.2. Computing Treatment Effects

One goal of the exogenous product removals is to determine how product-level sales respond to changes in availability. Let q_{jt} denote the sales of product j in machine-week t, superscript 1 denote sales when a focal product(s) is removed, and superscript 0 denote sales when a focal product(s) is available. Let the set of available products be A, and let F be the set of products we remove. Thus, $Q_t^1 = \sum_{j\in A\setminus F} q_{jt}^1$ and $Q_s^0 = \sum_{j\in A} q_{js}^0$ are the overall sales during treatment week t, and control week s respectively, and $q_{fs}^0 = \sum_{j\in F} q_{js}^0$ is the sales of the removed products during control week s. Our goal is to compute $\Delta q_{jt} = q_{jt}^1 - E[q_{jt}^0]$, the treatment effect of removing products(s) F on the sales of product j.

There are two challenges in implementing the removals and interpreting the data generated by them. The first challenge is that there is a large amount of variation in overall sales at the weekly level, independent of our exogenous removals. For example, a law firm may have a large case going to trial in a given month, and vend levels will increase at the firm during that period. In our particular setting, many of the product removals were done during the summer of 2007, which was a high-point in demand at these sites, most likely due to macroeconomic conditions. In this case, using a simple measure like previous weeks' sales, or overall average sales for $E[q_{jt}^0]$ could result in unreasonable treatment effects, such as sales increasing due to product removals, or sales decreasing by more than the sales of the focal products.

In order to deal with this challenge, we impose two simple restrictions based on consumer theory. Our first restriction is that our experimental product removals should not increase overall demand, so that $Q_t^0 - Q_s^1 \geq 0$ for treatment week t and control week s. Our second restriction is that the product removal(s) should not reduce overall demand by more than the sales of the products we removed, or $Q_t^0 - Q_s^1 \leq q_{fs}^0$. This means we choose control weeks s that correspond to treatment week t as follows:

$$\{s : s \neq t, Q_t^0 - Q_s^1 \in [0, q_{fs}^0]\}. \qquad (2)$$

While this has the nice property that it imposes the restriction on our selection of control weeks that all products are weak substitutes, it has the disadvantage that it introduces the potential for selection bias. The bias results from the fact that weeks with unusually high sales of the focal product q_{fs}^0 are more likely to be included in our control. This bias would likely overstate the costs of the product removal, which would be problematic for our study.

We propose a slight modification of ((2)) which removes the bias. That is, we replace q^0_{fs} with $\widehat{q^0_{fs}} = E[q^0_{fs}|Q^0_s]$. An easy way to obtain the expectation is to run an OLS regression of q^0_{fs} on Q^0_s, at the machine level, and use the predicted value. This has the nice property that the error is orthogonal to Q^0_s, which ensures that our choice of weeks is unbiased.

The second challenge is that, although the experimental design is relatively clean, the product mix presented in a machine is not necessarily fixed across machines, or within a machine over long periods of time, because we rely on observational data for the control weeks. For example, manufacturers may change their product lines, or Mark Vend may change its stocking decisions over time. Thus, while our field experiment intends to isolate the treatment effect of removing Snickers, we might instead compute the treatment effect of removing Snickers jointly with Mark Vend changing pretzel suppliers.

To mitigate this issue, we restrict our set of potential control weeks to those at the same machine with similar product availability within the category of our experiment. In practice, two of our three treatments took place during weeks where 3 Musketeers and Reese's Peanut Butter Cups were unavailable, so we restrict our set of potential control weeks for those experiments to weeks where those products were also unavailable. We denote this condition as $A_s \approx A_t$.

We use our definition of control weeks s to compute the expected control sales that correspond to treatment week t as:

$$S_t = \{s : s \neq t, A_t \approx A_s, Q^0_t - Q^1_s \in [0, \hat{b_0} + \hat{b_1} Q^0_s]\}. \tag{3}$$

And for each treatment week t we can compute the treatment effect as

$$\Delta q_{jt} = q^1_{jt} - \frac{1}{\#S_t} \sum_{s \in S_t} q^0_{js}. \tag{4}$$

While this approach has the advantage that it generates substitution patterns consistent with consumer theory, it may be the case that for some treatment weeks t the set of possible control weeks $S_t = \{\emptyset\}$. Under this definition of the control, some treatment weeks constitute 'outliers' and are excluded from the analysis. Of the 1470 machine-experiment-week combinations, 991 of them have at least one corresponding control week, and at the machine-experiment level, 528 out of 634 have at least one corresponding control. Each included treatment week has an average of 24 corresponding control weeks, though this can vary considerably from treatment week to treatment week.[5]

Once we have constructed our restricted set of treatment weeks and the set of control weeks that corresponds to each, inference is fairly straightforward. We use ((4)) to construct a set of pseudo-observations for the difference, and employ a paired t-test.

B.3. Estimation Algorithm

Here we provide pseudocode of our entire procedure for calculating $\pi(a, e)$. The first and third algorithm need to be repeated for each bootstrapped draw from the asymptotic distribution of $(\hat{d}_j, \hat{\sigma})$.

The computational 'trick' is to re-normalize the choice probabilities in Algorithm 1 steps 1(c-e). The normalization implicitly conditions on the set of customers who would have made a purchase at some hypothetical machine containing a superset of products $A_0 = A_t \cup \{(H, H), (H, M), (M, M)\}$. This can be justified in stages: the first stage is a draw from a binomial distribution where a consumer arrives and either selects the outside good or is labeled a 'likely consumer.' Likely consumers then face a second stage described by our re-normalized multinomial distribution where they choose either an available product or choose the outside good with a much smaller probability than the overall demand model $s_0(A_t) - \widetilde{s_0}(A_0)$. This saves time because we don't need to simulate the arrivals of consumers who never make a purchase. If the outside good share were 90% this would represent an order of magnitude reduction in the state space we ultimately need to keep track of as well as the number of consumer arrivals we need to simulate. This also makes the choice of ξ_t largely irrelevant as it governs the market share of the outside good and that gets normalized away. A larger ξ_t still increases the substitution probability to the outside option after products stock out. We calibrate this to $\xi = \text{med}(\xi_t) \approx 0.75$.[6]

If we were to increase ξ_t, this would decrease the share of the outside good and increase sales for any fixed number of consumers. However, because in Algorithm 2 we also estimate the arrival rate of consumers $P(x + \Delta x_k | x)$ in the normalized state-space, what happens instead is that as ξ_t increases we estimate a slower arrival process so that P is chosen to match the average daily sales observed in the top quartile of all machines across the entire MarkVend enterprise. We could have worked with the entire distribution of all machines, but

[5] Weeks in which the other five treatments were run (for the salty-snack and cookie categories) are excluded from the set of potential control weeks.

[6] We use the median because the distribution is highly skewed. We have also tried $\xi = E[\xi_t] = 0$, which gives nearly identical results. The optimal policies change by at most one unit.

we focus on this top quartile because we believe those machines drive restocking decisions – many of the slower machines are restocked because the driver is already nearby. A separate question is: "What is the point of ξ_t in the model?" and the answer is that we incorporate ξ_t in order to get unbiased estimates of $\hat{d}_j$ and $\widehat{\sigma}$.

Algorithm 1 Simulate Payoffs

1. Simulate consumer purchases from a full vending machine under assortment a.
   (a) Set $\xi = \text{med}[\widehat{\xi_t}] \approx 0.75$.
   (b) Initialize inventory of 15 confections products per slot for $a \in \{(H,H),(H,M),(M,M)\}$ plus products listed in table 8 (at modal max inventory). Label this inventory/assortment $A_0(a)$.
   (c) Use observed random coefficients demand parameters $(\widehat{d_j}, \widehat{\sigma})$ and quadrature nodes (w_i, ν_i) to calculate outside good purchase probability at an unobserved machine containing a superset of all possible products: $\tilde{A} = A_0 \cup \{(H,H),(H,M),(M,M)\}$.

$$\overline{s_0} = \sum_{i=1}^{NS} w_i \frac{e^{-\xi}}{\exp^{-\xi} + \sum_{k \in \tilde{A}} e^{\hat{d}_k + \sum_l \widehat{\sigma}_l \nu_{il} x_{kl}}}$$

   (d) Use observed random coefficients demand parameters $(\widehat{d_j}, \widehat{\sigma})$ and quadrature nodes (w_i, ν_i) to calculate purchase probabilities of a single consumer for current inventory/assortment A_t:

$$s_j(A_s) = \sum_{i=1}^{NS} w_i \frac{e^{\hat{d}_j + \sum_l \widehat{\sigma}_l \nu_{il} x_{jl}}}{\exp^{-\xi} + \sum_{k \in A_s} e^{\hat{d}_k + \sum_l \widehat{\sigma}_l \nu_{il} x_{kl}}}$$

   (e) Draw a single consumer purchase as y_t^*, a $(J+1)$ vector with re-normalized outside good probability.

$$y_s^* \sim Multinom\left(\frac{s_j(A_s)}{1-\overline{s_0}}, s_0(A_s) - \overline{s_0}\right)$$

   (f) Update $A_{s+1} = A_s - y_s^*$ or $A_{s+1} = A_s$ if outside good is chosen.
   (g) Continue for $s = 1, \ldots, 800$ consumers or (until machine is empty $A_s = \emptyset$).
   (h) Repeat for $n = 1, \ldots N = 100\,000$ machines to construct $y_{n,s}$: a $(J+1)$ vector.
2. Smooth Expected Flow Payoffs
   (a) Load retail and wholesale prices for all products. Assume $mc = 0.15$ for all confections.
   (b) Compute the expected flow payoffs for each agent as a function of cumulative arrivals x:

$$u^R(x,a) = \frac{1}{N} \sum_{n=1}^{N} \sum_{s=1}^{x} y_{n,s}^* \cdot (p_r - w)$$
$$u^M(x,a) = \frac{1}{N} \sum_{n=1}^{N} \sum_{s=1}^{x} y_{n,s}^* \cdot I_{nt}[\text{Mars}] \cdot (w_m - mc)$$
$$u^H(x,a) = \frac{1}{N} \sum_{n=1}^{N} \sum_{s=1}^{x} y_{n,s}^* \cdot I_{nt}[\text{Hershey}] \cdot (w_m - mc)$$
$$u^C(x,a) = \frac{1}{N} \sum_{n=1}^{N} \sum_{s=1}^{x} \log\left(1 + \sum_{j \in A_s} \exp\left[\hat{\delta}_k + \widehat{\xi} + \sum_l \widehat{\sigma}_l \nu_{il} x_{kl}\right]\right)$$

   (c) Smooth the expected profits $\left(u^R(x,a), u^M(x,a), u^H(x,a), u^C(x,a)\right) \rightarrow \left(\hat{u}^R(x,a), \hat{u}^M(x,a), \hat{u}^H(x,a), \hat{u}^C(x,a)\right)$ using MATLAB `slmegine`. Verify/require monotonicity for (R,M,C) but not (H,N).

Algorithm 2 Estimate the Arrival Rate

1. As in Algorithm 1 (Part 2) construct an estimate of total sales as a function of 'likely consumers'

$$u^{\text{sales}}(x,a) = \frac{1}{N}\sum_{n=1}^{N}\sum_{s=1}^{x}\sum_{j=1}^{J} y^{*}_{n,s}$$

2. For each visit in the data, measure the total sales Q_t since the previous service visit and calculate the fewest number of elapsed consumers x required to realize Q_t sales:

$$\hat{x}_t = \{\min x : u^{\text{sales}}(x,a) > Q_t\}$$

3. Denote the number of elapsed (business) days since the previous service visit as $days_t$ and define $\Delta x_t = \left(\frac{\hat{x}_t}{days_t}\right)$ as the average number of consumer arrivals per day for each visit t.

4. Construct a nonparametric frequency estimator for Δx_t:

$$P(\Delta x_t) = \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\left[b_k < \Delta x_t \leq b_{k+1}\right]$$

Algorithm 3 Solve the Dynamic Programming Problem

There exists a monotone policy such that the agent re-stocks if $x \geq e$:

1. Assume a known discount factor β and a fixed cost $FC = 10$.
2. Given a guess of the optimal policy, we can compute the post-decision pay-off $\tilde{u}$:
$$\tilde{u}(x, a, e) = \begin{cases} 0 & \text{if } x < e \\ \hat{u}(x, a) - FC & \text{if } x \geq e. \end{cases}$$
3. Compute the post-decision transition matrix $\tilde{P}$ by replacing columns of P.
$$\tilde{P}(x, e) = \begin{cases} x + \Delta x & \text{if } x < e \\ \Delta x & \text{if } x \geq e. \end{cases}$$
4. This allows us to solve the value function at all states in a single step:
$$V(x, a, e) = (I - \beta \tilde{P}(e))^{-1} \tilde{u}(x, a, e).$$
5. Find the ergodic/stationary distribution of x under policy e as the vector $\Gamma(e)$ that solves:
$$\Gamma(e) = \Gamma(e)\tilde{P}(x, e) \quad \text{with} \sum \Gamma(e) = 1.$$
6. Compute long-run expected profits under the Markov Chain using the stationary distribution:
$$\pi(a, e) = \Gamma(e)V(x, a, e)$$
7. Repeat this exercise for all possible choices of (a, e) and all agents R, M, H, N, C. Enumerate over e to find the optimal policy for each agent(s) (NR, R, VI, IND, SOC).

Algorithm 4 Compute the Standard Errors

1. Draw $\hat{\theta}^b \sim N\left(\hat{\theta}^{MLE}, \sqrt{diag(V(\hat{\theta}^{MLE})}\right)$. We only need: $(\hat{d}_j, \hat{\sigma})$. Assume $\xi = 0.75$ as before.
2. Simulate consumer arrivals and payoffs using Algorithm 1 $\hat{u}(x, a, \hat{\theta}^b)$ for each agent.
3. Use the same estimated consumer arrival process/ transition matrix $\hat{P}$ from Algorithm 2.
4. Use same calibrated discount factor β and same calibrated restocking cost $FC = 10$ and solve the dynamic programming problem using Algorithm 3.
5. Use $\pi^*(a, e|\hat{\theta}^b)$ to calculate the optimal policies for different groups of agents $(e^{NR}, e^R, e^{VI}, e^{SOC})$ for every (a, e) pair.
6. Compute all of the profit differences $\Delta\pi^R, \Delta\pi^M, \Delta\pi^H$.
7. Repeat 1000 times and report the standard deviations.

In this procedure there are two sources of variation. The first is the variation introduced by the uncertainty in the simulated ML estimates of the demand parameters (as reported in table **??**). The second is the simulation variance introduced from our simulation procedure, because we use the average over 100,000 chains this is designed to be at most ±$2.

B.4. Consumer Surplus and Welfare Calculations

Our calculation of the expected consumer surplus of a particular assortment and effort policy (a, e) parallels our calculation of retailer profits. We simulate consumer arrivals over many chains, and compute the set of available products as a function of the initial assortment a and the number of consumers to arrive since the previous restocking visit x which we write $a(x)$. For each assortment $a(x)$ that a consumer faces, we can compute the logit inclusive value and average over our simulations, to obtain an estimate at each x:

$$CS^*(a, x|\theta) = \frac{1}{I_t} \sum_{i=1}^{I_t} \log \left(\sum_{j \in a(x^s)} \exp[\delta_j + \mu_{ij}(\theta)] \right)$$

The exogenous arrival rate, $P(x'|x)$, denotes the expected daily number of consumer arrivals (from x cumulative likely consumers today to x' cumulative likely consumers tomorrow). Using this arrival rate and a policy e, we obtain the post-decision transition rule $\tilde{P}(x, e)$ and evaluate the ergodic distribution of consumer surplus under policy e:

$$CS^*(a, e) = [I - \beta \tilde{P}(x, e)]^{-1} CS^*(a, x|\theta)$$

The remaining challenge is that $CS^*(a,e)$ relates to arbitrary units of consumer utility, rather than dollars. Recall our utility specification from Equation (1) in Section 4 of the main text, with $\theta = [\delta, \alpha, \sigma]$:

$$u_{ijt}(\theta) = \delta_j + \alpha p_{jt} + \xi_t + \sum_l \sigma_l \nu_{ilt} x_{jl} + \varepsilon_{ijt}$$

Without observable, within-product variation in price, $p_{jt} = p_j$, and α is not separately identified from the product fixed-effect δ_j. If α were identified, then we could simply write $CS(a,e) = \frac{1}{\alpha} CS^*(a,e)$. Instead, we can calibrate α given an own price elasticity:

$$\epsilon_{j,t} = \frac{p_{jt}}{s_{jt}} \cdot \frac{\partial s_{jt}}{\partial p_{jt}} = \frac{p_{jt}}{s_{jt}} \cdot \int \frac{\partial s_{ijt}}{\partial p_{jt}} f(\beta_i|\theta) d\,\beta_i = \alpha \cdot \underbrace{\frac{p_{jt}}{s_{jt}} \cdot \int (1 - s_{ij}(\delta, \beta_i)) \cdot s_{ij}(\delta, \beta_i) f(\beta_i|\theta) d\,\beta_i}_{\epsilon^*_{j,t}(\theta)}$$

The term $\epsilon^*_{j,t}$ does not depend directly on α once we have controlled for the fixed effect d_j. Thus, we can calibrate own-price elasticities. As is conventional in the literature, we work with the median own-price elasticity, $\bar{\epsilon}(\theta) = \text{median}_j(\epsilon^*_{j,t}(\theta))$, and recover α as $\alpha = |\frac{\epsilon}{\bar{\epsilon}(\theta)}|$. We then calculate α at three different values of ϵ: $\epsilon \in \{-1, -2, -4\}$.

As is well known, α has an alternative interpretation in the social planner's problem as the planner's weight on consumer surplus:

$$SS(a,e) = PS(a,e) + \frac{\gamma}{|\alpha|} CS^*(a,e)$$

The social planner's problem is equivalent in the following cases: (1) the median own-price elasticity is $\epsilon = -2$ and $\gamma = 1$; (2) the median own-price elasticity is $\epsilon = -4$ and the planner puts twice as much weight on consumer surplus $\gamma = 2$; (3) the median own-price elasticity is $\epsilon = -1$ and the planner puts half as much weight on consumer surplus $\gamma = \frac{1}{2}$.

C. Robustness Checks

For each of our robustness checks we change the parameters of the dynamic decision problem and see if it changes the welfare implications of the AUD contract. To summarize these results, we compare our alternative specifications to Table 13 from the main text. This allows us to compare both foreclosure and efficiency effects at the same time. We focus on some key outcomes, the first is the sign of the change in producer and consumer surplus for transitions between $(H, H) \rightarrow (M, M)$ under different effort levels and from $(H, M) \rightarrow (M, M)$. In nearly all of the robustness test we find results qualitatively similar to those in the main text. First, both consumers and producers are better off under the (H, M) assortment than the (M, M) assortment. Second, the overall impact on consumers is sometimes ambiguous as they can be compensated for an inferior assortment with a higher effort level under e^{VI}. As in the main text this depends on the retailer setting a lower effort level e^{NR} under the (H, M) assortment. Third, Hershey would have to set a very low wholesale price (often below our assumed 15 cent marginal cost) in order to avoid being foreclosed. Similarly, this implies that Mars could only modestly reduce the generosity of the rebate (by 4-6%) without Hershey being able to respond and avoid foreclosure.

We consider a broad array of alternatives: changing the arrival rate of consumers; setting the marginal cost to zero and maximizing potential efficiencies; increasing or decreasing the fixed cost of restocking; and having the retailer place some weight on consumer surplus when making decisions.

C.1. Arrival Rate: Details and Robustness

We estimate the arrival rate $P(\Delta x_t)$ by grouping machines across the entire MarkVend enterprise into quartiles based on average daily sales for the entire sample. Our main specification focuses on the top quartile of machines by this metric. As a robustness test, we also consider the next 50% (25th to 75th percentile machines). For each machine-visit we calculate the average daily sales and the total sales when the machine was restocked. The first metric can be used to estimate $P(\Delta x_t)$ while the second metric can be thought of as an empirical estimate of the policy function $e(\cdot)$. Neither of these are strictly correct because some consumers arrive at the machine and elect to purchase the outside option. However, in our normalized state space x_t represents the cumulative number of consumer arrivals since our last restocking event, *who would have purchased at a full machine.* Thus the only gap arises from consumers who would have purchased at a full machine but do not purchase because of stockouts. For $x_t \leq 300$ consumer arrivals this implies an adjustment of $\leq 10\%$ between the policy in the space of realized sales and consumer arrivals in the model.

In Figure A5 we replicate Figure 4 from the main text above and below include the middle 50% of machines. We see that the arrival rate is substantially lower for the middle 50% of machines (15.4 per day) than for the top 25% of machines (37.6 per day) as we might expect. We also see that the empirical distribution of restocking policies for these machines is lower (mean of $e \approx 130$ versus mean of $e \approx 80$). This does not imply that MarkVend services less popular machines *more frequently* but rather they service less popular machines *after fewer consumer arrivals*; the confound is the lower arrival rate at these machines. A likely story is that these machines have lower fixed costs to service (perhaps because the driver is already on-site servicing a nearby machine, or because it takes less time to restock fewer products). This is part of the reason we chose to focus on machines with above average consumer arrival rates, because we believe those are more likely to drive MarkVend's stocking decisions.

An important question is whether or results are sensitive to the arrival rate of consumers. We reproduce the 'net effects' table (Table 13 from the main text) as Table A1 below. We find that all of the qualitative results are the same: the rebate can be used to foreclose the rival even though (H, M) generates more producer surplus than (M, M). Overall welfare impacts are the same as in the main text. The (H, M) assortment maximizes producer surplus and consumer surplus. It is possible that consumers receive sufficient benefits in moving from e^{NR} to e^{VI} to compensate them for the inferior assortment (M, M), though e^R does not provide sufficient compensation.

Table A1: Net Effect of Efficiency and Foreclosure (Middle 50% of Machines)

from to (M,M) and:	(H,M) and e^{NR} e^R	 e^{VI}	 e^{SOC}	(H,H) and e^{NR} e^R	 e^{VI}	 e^{SOC}
$\Delta\pi^R$	-304	-350	-570	-646	-692	-912
$\Delta\pi^M$	1094	1174	1285	2362	2442	2553
$\Delta\pi^H$	-908	-908	-908	-1518	-1518	-1518
$\Delta\pi^N$	-4	-6	-7	-22	-24	-25
ΔPS	-123	-91	-201	176	207	97
$\Delta CS(\epsilon = -2)$	-22	153	406	222	398	650
ΔSS	-145	62	204	398	605	747
$\lambda\pi^M$	2321	2339	2364	2321	2339	2364
w_h to avoid foreclosure	-18.87	-18	-12.03	12.12	12.64	16.22
Reduced λ (Percent)	47.76	46.18	37.44	6.78	5.51	-2.81

Notes: Consumer Surplus calibrates α to median own price elasticity of $\epsilon = -2$. Calibration only affects the scale of consumer surplus calculations, not the ranking of various options. For more details see Appendix B.4.

We tried alternative assumptions on the arrival rate by either doubling or halving the

Figure A5: Observed Policies and Arrival Rates

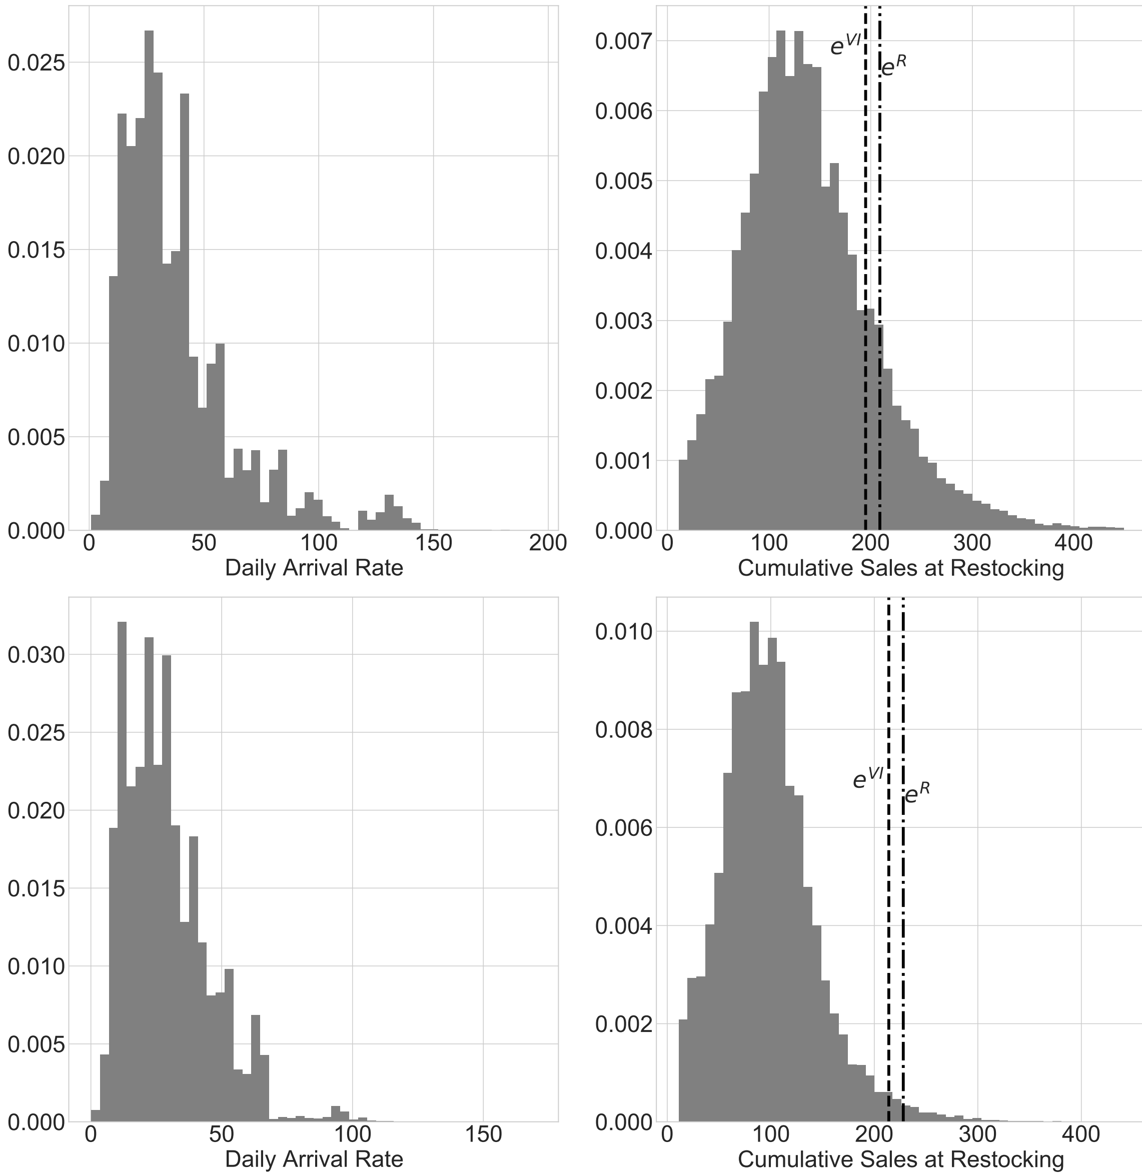

Notes: Top row reports daily arrival rate for top 25% of machines at MarkVend's overall enterprise. Bottom row reports daily arrival rate for middle 50% of MarkVend's machines. These are used to estimate $f(\Delta x_t)$. Right column reports cumulative sales at restocking as well as calculated optimal policies from the model. Policies and cumulative sales are in the same units except for 'sales' of the outside good.

rate at which customers arrive. Though we don't report those results here, we didn't find a substantial effect on anything other than the absolute magnitude of profits.

C.2. Robustness to Alternative Marginal Costs

We reproduce the 'net effects' as Table A2 where we set the marginal cost of production equal to zero. The main difference is that manufacturer profits are larger in all scenarios. The gap between the retailer optimal policy e^R and the vertically integrated e^{VI} or socially optimal e^{SOC} policy becomes larger. This can be viewed as a way to obtain an 'upper bound' on potential efficiencies as now production is costless. We find that all of the qualitative results and signs of point estimates are the same: the rebate can be used to foreclose the rival even though (H, M) generates more producer surplus than (M, M). Hershey's countermeasures are similar to those we calculated in the main text. It would have to cut its wholesale price below 15 cents to avoid foreclosure under both e^R and e^{VI}. Likewise, Mars could not reduce the rebate by much and still foreclose Hershey: only 4% at the vertically-integrated effort level and 6.6% at e^R.

Overall, welfare impacts are the same as in the main text. The (H, M) assortment maximizes producer surplus and consumer surplus. It is possible that consumers receive sufficient benefit in moving from e^{NR} to e^{VI} to compensate them for the inferior assortment (M, M), though e^R, and does not provide sufficient compensation.

Table A2: Net Effect of Efficiency and Foreclosure ($MC = 0$)

from to (M,M) and:	(H,M) and e^{NR} e^R	 e^{VI}	 e^{SOC}	(H,H) and e^{NR} e^R	 e^{VI}	 e^{SOC}
$\Delta\pi^R$	-733	-930	-1467	-1548	-1746	-2282
$\Delta\pi^M$	3641	4009	4320	7868	8236	8547
$\Delta\pi^H$	-3361	-3361	-3361	-5614	-5614	-5614
$\Delta\pi^N$	-16	-23	-25	-82	-89	-91
ΔPS	-469	-306	-533	624	787	560
$\Delta CS(\epsilon = -2)$	-55	534	1047	534	1122	1636
ΔSS	-524	228	514	1157	1909	2195
$\lambda\pi^M$	5546	5605	5655	5546	5605	5655
w_h to avoid foreclosure	-18.48	-16.72	-10.53	12.31	13.36	17.06
Reduced λ (Percent)	47.45	44.48	35.49	6.38	3.84	-4.79

Notes: Consumer Surplus calibrates α to median own price elasticity of $\epsilon = -2$. Calibration only affects the scale of consumer surplus calculations, not the ranking of various options. For more details see Appendix B.4.

C.3. Robustness to Alternative Fixed Costs

We reproduce the 'net effects' from the text as Tables A3 and A4 below. The main response to the fixed cost is that potential efficiency effects are smaller when the fixed costs are smaller and larger when the fixed costs are greater. Higher fixed costs reduce both the profits and the effort level of the retailer.

We find that all of the qualitative results are the same: the rebate can be used to foreclose the rival even though (H, M) generates more producer surplus than (M, M). The point estimates all have the same sign as those in Table 13, though for $FC = 15$ the sign flips on ΔCS when moving from (H, M) and e^{NR} to e^R and (M, M). Thus even at the vertically integrated effort level, it is impossible to compensate consumers for the inferior assortment.

The effect on rival countermeasures are similar: at the lower fixed cost Hershey would need to reduce prices even more than in the main text to avoid foreclosure; while Mars could reduce the generosity of the rebate slightly more (around 7%); at the higher fixed cost Hershey would need to reduce prices less than in the main text to avoid foreclosure; while Mars could reduce the generosity of the rebate slightly less (around 5%).

Table A3: Net Effect of Efficiency and Foreclosure ($FC = 5$)

from	(H,M) and e^{NR}			(H,H) and e^{NR}		
to (M,M) and:	e^R	e^{VI}	e^{SOC}	e^R	e^{VI}	e^{SOC}
$\Delta\pi^R$	-656	-698	-950	-1622	-1664	-1915
$\Delta\pi^M$	2549	2621	2741	5542	5614	5734
$\Delta\pi^H$	-2168	-2168	-2168	-3631	-3631	-3631
$\Delta\pi^N$	-9	-8	-3	-55	-54	-50
ΔPS	-285	-253	-381	234	266	138
$\Delta CS(\epsilon = -2)$	-183	-9	289	426	600	897
ΔSS	-468	-262	-92	660	865	1035
$\lambda\pi^M$	5670	5686	5712	5670	5686	5712
w_h to avoid foreclosure	-21.41	-21.08	-18.2	11.81	12.01	13.73
Reduced λ (Percent)	50.18	49.59	45.41	7.36	6.89	2.92

Notes: Consumer Surplus calibrates α to median own price elasticity of $\epsilon = -2$. Calibration only affects the scale of consumer surplus calculations, not the ranking of various options. For more details see Appendix B.4.

Table A4: Net Effect of Efficiency and Foreclosure ($FC = 15$)

from	(H,M) and e^{NR}			(H,H) and e^{NR}		
to (M,M) and:	e^R	e^{VI}	e^{SOC}	e^R	e^{VI}	e^{SOC}
$\Delta\pi^R$	-792	-971	-1788	-1509	-1688	-2505
$\Delta\pi^M$	2679	2992	3414	5777	6090	6512
$\Delta\pi^H$	-2202	-2202	-2202	-3673	-3673	-3673
$\Delta\pi^N$	-12	-21	-30	-54	-63	-73
ΔPS	-328	-203	-607	541	666	261
$\Delta CS(\epsilon = -2)$	24	702	1630	606	1284	2212
ΔSS	-303	500	1023	1147	1950	2474
$\lambda\pi^M$	5426	5496	5590	5426	5496	5590
w_h to avoid foreclosure	-15.64	-14.26	-5.15	13.15	13.98	19.45
Reduced λ (Percent)	44.82	42.25	28.61	4.5	2.45	-10.53

Notes: Consumer Surplus calibrates α to median own price elasticity of $\epsilon = -2$. Calibration only affects the scale of consumer surplus calculations, not the ranking of various options. For more details see Appendix B.4.

C.4. Joint Retailer-Consumer Surplus

We also allow the retailer to optimize the joint surplus of the retailer and the consumer. This may be an important consideration if providing good service to the consumer is an important aspect of how our retail operator competes with other vending operators for contracts with retail locations. It may also help explain why our retailer provides an extremely high frequency of service visits (beyond what we can justify with an optimal stocking model). We find that for $\epsilon = -1$ and $\gamma = 3$ so that $\frac{\gamma}{\alpha} = 6$, we are able to produce an effort policy which matches the mean of the observed distribution of retailer effort in Figure 4 of $e \approx 130$.

Table A5 reports the optimal effort policies of a joint Retailer-Consumer entity. By placing a large weight on consumer surplus, the retailer substantially increases its effort under all assortments. Also, because the resulting effort level is so high the potential efficiency effects of the rebate are highly limited and the gap between the effort set by the retailer and e^{VI} is quite small.

Table A5: Optimal Effort Policies: Restock after how many customers?

	(M,H)	(H,H)	(M,M)	(M,H)	(H,H)	(M,M)
	Effort Policy			% Change from e^{NR}		
e^{NR}	130	130	130	0.00	0.00	0.00
e^{R}	130	130	130	0.00	0.00	0.00
e^{VI}	130	130	130	0.00	0.00	0.00
e^{IND}	130	130	130	0.00	0.00	0.00
e^{SOC}	172	168	171	164.62	167.69	165.38
e^{SOC1}	157	154	156	176.15	178.46	176.92
e^{SOC4}	183	178	181	156.15	160.00	157.69

Notes: Reported for retailer who places weight $\frac{\gamma}{\alpha} = 6$ on consumer surplus. For further details, see Appendix B.4. The width of the 95% CI is at most one unit.

The potential gains are much smaller than they are in the case where the retailer does not take consumer surplus into account. For all elasticities, the potential change in the restocking frequency is now less than 5%. Likewise, the maximum change in social surplus is less than \$75 for all elasticities and assortments. Once the retailer internalizes the effect of effort on consumers, there is little to be gained from internalizing the same effort effect on the upstream manufacturer. The retailer-consumer pair exerts more effort than the vertically integrated retailer-Mars pair in our base scenario.

Though it is likely in practice that MarkVend at least partially considers consumer surplus when choosing its effort level, our base scenario ignores this possibility. Incorporating consumer surplus in the retailer's effort decision drastically reduces potential efficiency effects of the rebate contract. Ultimately, we are interested in whether an efficiency effect might outweigh potential foreclosure effects, and we design our baseline estimates to be an ‘upper bound’ on such effects.

Table A6: Net Effect of Efficiency and Foreclosure

from to (M,M) and:	(H,M) and e^{NR} e^{R}	 e^{VI}	 e^{SOC}	(H,H) and e^{NR} e^{R}	 e^{VI}	 e^{SOC}
$\Delta\pi^{R}$	-616	-616	1624	-1679	-1679	561
$\Delta\pi^{M}$	2507	2507	2199	5514	5514	5207
$\Delta\pi^{H}$	-2176	-2176	-2176	-3641	-3641	-3641
$\Delta\pi^{N}$	-11	-11	-18	-58	-58	-65
ΔPS	-296	-296	1630	137	137	2062
$\Delta CS(\epsilon=-2)$	-275	-275	-1021	422	422	-324
ΔSS	-571	-571	609	559	559	1738
$\lambda\pi^{M}$	5718	5718	5649	5718	5718	5649
w_h to avoid foreclosure	-22.31	-22.31	-50.01	11.97	11.97	-4.59
Reduced λ (Percent)	51.17	51.17	90.23	6.96	6.96	45.49

Notes: Reported for retailer who places weight $\frac{\gamma}{\alpha}=6$ on consumer surplus. For more details see Appendix B.4.

D. Full $\pi(a,e)$ Tables

We compute $\pi(a,e)$ for every agent and 15 assortments. We report only the most relevant assortments and effort levels below. Note that $\pi(a,e)$ denotes the present discounted value of profits from a single machine in the top quartile of the MarkVend enterprise. We cannot report exact profits at the enterprise level but it is safe to assume they are orders of magnitude larger. First column reports policy type and value in parentheses.

Policy	π^R	$\lambda\pi^M$	π^M	π^H	π^N	$\pi^R + \pi^M$	PS	CS
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
e^{NR}(217)	94,733	4,964	22,363	2,181	2,149	117,095	121,425	65,491
e^{R}(211)	94,723	4,985	22,454	2,179	2,146	117,177	121,502	65,685
e^{VI}(197)	94,612	5,028	22,648	2,173	2,143	117,260	121,576	66,105
e^{IND}(197)	94,612	5,028	22,648	2,173	2,143	117,260	121,576	66,105
e^{SOC}(172)	94,060	5,091	22,934	2,168	2,141	116,994	121,303	66,738
e^{SOC1}(157)	93,469	5,121	23,068	2,169	2,142	116,536	120,848	67,048
e^{SOC4}(183)	94,363	5,066	22,818	2,170	2,141	117,181	121,492	66,478
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
$e^{Pre2008}$(137)	92,296	5,152	23,207	2,174	2,147	115,503	119,824	67,387
$e^{Post2008}$(144)	92,768	5,142	23,162	2,172	2,145	115,931	120,247	67,276
e^{NR}(212)	95,548	4,288	19,316	3,644	2,192	114,864	120,700	64,902
e^{R}(206)	95,537	4,310	19,415	3,641	2,190	114,952	120,783	65,095
e^{VI}(191)	95,407	4,360	19,642	3,634	2,187	115,048	120,869	65,539
e^{IND}(191)	95,407	4,360	19,642	3,634	2,187	115,048	120,869	65,539
e^{SOC}(168)	94,876	4,424	19,926	3,630	2,187	114,802	120,619	66,111
(M,M) Assortment: Three Musketeers and Milkyway								
e^{SOC1}(154)	94,316	4,454	20,063	3,632	2,189	114,379	120,200	66,398
e^{SOC4}(178)	95,161	4,398	19,812	3,631	2,186	114,972	120,789	65,878
$e^{Pre2008}$(137)	93,339	4,483	20,194	3,637	2,194	113,533	119,364	66,688
$e^{Post2008}$(144)	93,791	4,472	20,144	3,635	2,192	113,934	119,761	66,574
e^{NR}(217)	94,005	5,521	24,867	0	2,141	118,872	121,013	65,173
e^{R}(211)	94,005	5,541	24,958	0	2,139	118,962	121,101	65,371
e^{VI}(197)	93,915	5,584	25,152	0	2,135	119,067	121,201	65,801
e^{IND}(197)	93,915	5,584	25,152	0	2,135	119,067	121,201	65,801
e^{SOC}(172)	93,397	5,647	25,438	0	2,132	118,835	120,967	66,448
e^{SOC1}(157)	92,825	5,677	25,572	0	2,133	118,397	120,530	66,765
e^{SOC4}(183)	93,686	5,621	25,322	0	2,132	119,008	121,141	66,182
$e^{Pre2008}$(137)	91,673	5,708	25,713	0	2,137	117,387	119,523	67,111
$e^{Post2008}$(144)	92,139	5,698	25,668	0	2,135	117,807	119,942	66,998

Table A7: Simulated Profits for Main Specification

Notes: Profit numbers represent the long-run expected profit from a top quartile machine. Rebate payments are assumed to only be paid under an (M, M) assortment; rebate payments are assumed to not be paid to the retailer. The socially-optimal assortment is (H, M). First column reports policy type and value in parenthesis. $FC = 10, MC = 0.15$.

Policy	π^R	$\lambda\pi^M$	π^M	π^H	π^N	$\pi^R+\pi^M$	PS	CS
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
$e^{NR}(217)$	94,733	4,964	31,024	3,361	3,260	125,756	132,377	65,491
$e^{R}(211)$	94,723	4,984	31,150	3,356	3,257	125,873	132,486	65,685
$e^{VI}(191)$	94,523	5,044	31,525	3,345	3,250	126,048	132,643	66,271
$e^{IND}(192)$	94,539	5,041	31,508	3,345	3,250	126,048	132,643	66,244
$e^{SOC}(169)$	93,959	5,097	31,856	3,340	3,248	125,815	132,403	66,804
$e^{SOC1}(155)$	93,372	5,124	32,024	3,342	3,251	125,396	131,989	67,086
$e^{SOC4}(179)$	94,265	5,075	31,716	3,341	3,248	125,981	132,570	66,576
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
$e^{Pre2008}(137)$	92,296	5,151	32,195	3,348	3,258	124,491	131,098	67,387
$e^{Post2008}(144)$	92,768	5,141	32,133	3,345	3,255	124,901	131,501	67,276
$e^{NR}(212)$	95,548	4,287	26,797	5,614	3,326	122,344	131,284	64,902
$e^{R}(206)$	95,537	4,310	26,935	5,609	3,323	122,472	131,403	65,095
$e^{VI}(185)$	95,310	4,378	27,362	5,595	3,317	122,672	131,584	65,700
$e^{IND}(186)$	95,328	4,375	27,343	5,596	3,317	122,671	131,585	65,674
$e^{SOC}(165)$	94,772	4,430	27,687	5,593	3,318	122,459	131,370	66,176
(M,M) Assortment: Three Musketeers and Milkyway								
$e^{SOC1}(152)$	94,219	4,457	27,857	5,596	3,322	122,076	130,994	66,435
$e^{SOC4}(174)$	95,058	4,408	27,550	5,593	3,317	122,608	131,518	65,974
$e^{Pre2008}(137)$	93,339	4,482	28,016	5,604	3,329	121,354	130,287	66,688
$e^{Post2008}(144)$	93,791	4,471	27,945	5,599	3,325	121,736	130,660	66,574
$e^{NR}(217)$	94,005	5,520	34,498	0	3,248	128,503	131,751	65,173
$e^{R}(211)$	94,005	5,540	34,624	0	3,245	128,629	131,873	65,371
$e^{VI}(191)$	93,835	5,600	34,999	0	3,237	128,833	132,070	65,970
$e^{IND}(192)$	93,850	5,597	34,982	0	3,237	128,831	132,069	65,943
$e^{SOC}(169)$	93,300	5,653	35,330	0	3,234	128,631	131,865	66,516
$e^{SOC1}(155)$	92,731	5,680	35,499	0	3,236	128,230	131,466	66,804
$e^{SOC4}(179)$	93,593	5,630	35,190	0	3,235	128,783	132,018	66,282
$e^{Pre2008}(137)$	91,673	5,708	35,672	0	3,242	127,345	130,587	67,111
$e^{Post2008}(144)$	92,139	5,697	35,609	0	3,239	127,748	130,987	66,998

Table A8: Simulated Profits for $MC = 0$ Specification

Notes: Profit numbers represent the long-run expected profit from a top quartile machine. Rebate payments are assumed to only be paid under an (M, M) assortment; rebate payments are assumed to not be paid to the retailer. The socially-optimal assortment is (H, M). First column reports policy type and value in parenthesis. $FC = 10, MC = 0$.

Policy	π^R	$\lambda\pi^M$	π^M	π^H	π^N	$\pi^R + \pi^M$	PS	CS
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
e^{NR}(236)	39,596	2,078	9,362	908	895	48,958	50,761	27,388
e^{R}(231)	39,593	2,086	9,394	907	894	48,987	50,788	27,456
e^{VI}(216)	39,544	2,105	9,481	905	892	49,025	50,822	27,643
e^{IND}(217)	39,549	2,104	9,476	905	892	49,025	50,822	27,631
e^{SOC}(192)	39,325	2,130	9,595	903	891	48,920	50,714	27,894
e^{SOC1}(177)	39,084	2,142	9,650	903	892	48,733	50,528	28,022
e^{SOC4}(202)	39,439	2,120	9,552	903	891	48,991	50,785	27,797
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
$e^{Pre2008}$(137)	37,885	2,163	9,743	909	898	47,629	49,436	28,265
$e^{Post2008}$(144)	38,168	2,160	9,731	908	897	47,900	49,704	28,231
e^{NR}(231)	39,938	1,797	8,095	1,518	913	48,032	50,463	27,143
e^{R}(225)	39,933	1,806	8,136	1,516	912	48,069	50,498	27,224
e^{VI}(210)	39,878	1,827	8,231	1,513	911	48,109	50,533	27,408
e^{IND}(211)	39,884	1,826	8,225	1,513	911	48,109	50,533	27,396
e^{SOC}(188)	39,670	1,852	8,344	1,512	910	48,014	50,436	27,634
(M,M) Assortment: Three Musketeers and Milkyway								
e^{SOC1}(174)	39,442	1,865	8,400	1,513	912	47,842	50,266	27,752
e^{SOC4}(197)	39,775	1,843	8,302	1,512	910	48,077	50,499	27,548
$e^{Pre2008}$(137)	38,348	1,886	8,496	1,521	918	46,844	49,283	27,977
$e^{Post2008}$(144)	38,623	1,883	8,482	1,519	917	47,105	49,541	27,943
e^{NR}(236)	39,294	2,310	10,406	0	892	49,701	50,592	27,256
e^{R}(231)	39,295	2,317	10,438	0	891	49,733	50,624	27,325
e^{VI}(216)	39,255	2,337	10,525	0	889	49,781	50,670	27,517
e^{IND}(217)	39,260	2,335	10,520	0	889	49,780	50,669	27,505
e^{SOC}(192)	39,051	2,362	10,639	0	888	49,690	50,578	27,775
e^{SOC1}(177)	38,818	2,374	10,694	0	888	49,512	50,400	27,905
e^{SOC4}(202)	39,160	2,352	10,596	0	888	49,755	50,643	27,675
$e^{Pre2008}$(137)	37,635	2,395	10,790	0	894	48,425	49,319	28,152
$e^{Post2008}$(144)	37,916	2,393	10,777	0	892	48,694	49,586	28,118

Table A9: Simulated Profits for Middle 50% machines

Notes: Profit numbers represent the long-run expected profit from a 25-75 percentile machine. Rebate payments are assumed to only be paid under an (M, M) assortment; rebate payments are assumed to not be paid to the retailer. The socially-optimal assortment is (H, M). First column reports policy type and value in parenthesis. $FC = 10, MC = 0.15$.

Policy	π^R	$\lambda\pi^M$	π^M	π^H	π^N	$\pi^R + \pi^M$	PS	CS
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
$e^{NR}(166)$	101,096	5,104	22,991	2,168	2,141	124,087	128,396	66,869
$e^{R}(163)$	101,093	5,110	23,017	2,168	2,141	124,111	128,420	66,931
$e^{VI}(154)$	101,052	5,126	23,091	2,169	2,143	124,144	128,456	67,105
$e^{IND}(153)$	101,045	5,128	23,099	2,170	2,143	124,144	128,456	67,123
$e^{SOC}(135)$	100,787	5,155	23,219	2,174	2,148	124,007	128,329	67,418
$e^{SOC1}(130)$	100,669	5,161	23,248	2,176	2,150	123,917	128,243	67,492
$e^{SOC4}(143)$	100,932	5,144	23,169	2,172	2,145	124,101	128,419	67,293
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
$e^{Pre2008}(137)$	100,829	5,152	23,207	2,174	2,147	124,036	128,357	67,387
$e^{Post2008}(144)$	100,947	5,142	23,162	2,172	2,145	124,109	128,426	67,276
$e^{NR}(161)$	102,062	4,439	19,998	3,631	2,188	122,059	127,877	66,260
$e^{R}(158)$	102,059	4,446	20,026	3,631	2,188	122,085	127,904	66,320
$e^{VI}(148)$	102,008	4,465	20,113	3,633	2,190	122,121	127,945	66,507
$e^{IND}(148)$	102,008	4,465	20,113	3,633	2,190	122,121	127,945	66,507
$e^{SOC}(131)$	101,755	4,492	20,234	3,640	2,196	121,989	127,826	66,780
(M,M) Assortment: Three Musketeers and Milkyway								
$e^{SOC1}(130)$	101,732	4,493	20,241	3,641	2,197	121,973	127,810	66,795
$e^{SOC4}(138)$	101,888	4,482	20,187	3,637	2,194	122,075	127,906	66,673
$e^{Pre2008}(137)$	101,872	4,483	20,194	3,637	2,194	122,066	127,897	66,688
$e^{Post2008}(144)$	101,969	4,472	20,144	3,635	2,192	122,113	127,939	66,574
$e^{NR}(166)$	100,442	5,660	25,495	0	2,132	125,937	128,068	66,581
$e^{R}(163)$	100,442	5,666	25,522	0	2,132	125,964	128,096	66,645
$e^{VI}(154)$	100,412	5,682	25,596	0	2,133	126,008	128,141	66,822
$e^{IND}(153)$	100,405	5,684	25,604	0	2,133	126,009	128,143	66,841
$e^{SOC}(135)$	100,167	5,711	25,726	0	2,137	125,893	128,030	67,142
$e^{SOC1}(130)$	100,053	5,718	25,755	0	2,139	125,808	127,947	67,217
$e^{SOC4}(143)$	100,304	5,700	25,675	0	2,135	125,978	128,114	67,015
$e^{Pre2008}(137)$	100,206	5,708	25,713	0	2,137	125,920	128,057	67,111
$e^{Post2008}(144)$	100,317	5,698	25,668	0	2,135	125,985	128,120	66,998

Table A10: Simulated Profits for $FC = 5$

Notes: Profit numbers represent the long-run expected profit from a top quartile machine with $FC = 5$. Rebate payments are assumed to only be paid under an (M, M) assortment; rebate payments are assumed to not be paid to the retailer. The socially-optimal assortment is (H, M). First column reports policy type and value in parenthesis. $FC = 5, MC = 0.15$.

Policy	π^R	$\lambda\pi^M$	π^M	π^H	π^N	$\pi^R+\pi^M$	PS	CS
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
e^{NR}(251)	89,383	4,832	21,764	2,202	2,165	111,147	115,514	64,228
e^{R}(244)	89,369	4,861	21,898	2,198	2,161	111,267	115,626	64,510
e^{VI}(226)	89,185	4,932	22,218	2,186	2,152	111,403	115,741	65,183
e^{IND}(227)	89,201	4,929	22,201	2,187	2,153	111,402	115,742	65,148
e^{SOC}(197)	88,372	5,028	22,648	2,173	2,143	111,020	115,336	66,105
e^{SOC1}(179)	87,477	5,075	22,862	2,169	2,141	110,339	114,649	66,576
e^{SOC4}(209)	88,794	4,991	22,483	2,178	2,146	111,277	115,600	65,748
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
$e^{Pre2008}$(137)	83,763	5,152	23,207	2,174	2,147	106,970	111,291	67,387
$e^{Post2008}$(144)	84,590	5,142	23,162	2,172	2,145	107,752	112,069	67,276
e^{NR}(246)	90,100	4,144	18,666	3,673	2,207	108,766	114,646	63,646
e^{R}(239)	90,084	4,176	18,811	3,666	2,203	108,894	114,764	63,927
e^{VI}(220)	89,871	4,257	19,176	3,650	2,195	109,047	114,891	64,632
e^{IND}(221)	89,889	4,253	19,158	3,651	2,195	109,046	114,892	64,597
e^{SOC}(193)	89,080	4,354	19,613	3,635	2,187	108,693	114,515	65,483
(M,M) Assortment: Three Musketeers and Milkyway								
e^{SOC1}(176)	88,223	4,404	19,836	3,631	2,186	108,059	113,875	65,926
e^{SOC4}(205)	89,510	4,314	19,431	3,640	2,190	108,941	114,771	65,126
$e^{Pre2008}$(137)	84,805	4,483	20,194	3,637	2,194	105,000	110,831	66,688
$e^{Post2008}$(144)	85,612	4,472	20,144	3,635	2,192	105,756	111,582	66,574
e^{NR}(251)	88,599	5,389	24,273	0	2,157	112,872	115,030	63,885
e^{R}(244)	88,597	5,418	24,406	0	2,154	113,003	115,157	64,172
e^{VI}(226)	88,443	5,488	24,723	0	2,145	113,166	115,310	64,859
e^{IND}(227)	88,457	5,485	24,706	0	2,145	113,164	115,309	64,823
e^{SOC}(197)	87,675	5,584	25,152	0	2,135	112,827	114,962	65,801
e^{SOC1}(179)	86,806	5,631	25,366	0	2,132	112,172	114,304	66,282
e^{SOC4}(209)	88,078	5,547	24,987	0	2,138	113,066	115,204	65,436
$e^{Pre2008}$(137)	83,140	5,708	25,713	0	2,137	108,854	110,990	67,111
$e^{Post2008}$(144)	83,960	5,698	25,668	0	2,135	109,628	111,763	66,998

Table A11: Simulated Profits for $FC = 15$

Notes: Profit numbers represent the long-run expected profit from a top quartile machine with $FC = 15$. Rebate payments are assumed to only be paid under an (M, M) assortment; rebate payments are assumed to not be paid to the retailer. The socially-optimal assortment is (H, M). First column reports policy type and value in parenthesis. $FC = 15, MC = 0.15$.

Policy	π^R	$\lambda\pi^M$	π^M	π^H	π^N	$\pi^R+\pi^M$	PS	CS
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
$e^{NR}(130)$	91,742	5,161	23,248	2,176	2,150	114,990	119,316	67,492
$e^{R}(130)$	91,742	5,161	23,248	2,176	2,150	114,990	119,316	67,492
$e^{VI}(130)$	91,742	5,161	23,248	2,176	2,150	114,990	119,316	67,492
$e^{IND}(130)$	91,742	5,161	23,248	2,176	2,150	114,990	119,316	67,492
$e^{SOC}(172)$	94,060	5,091	22,934	2,168	2,141	116,994	121,303	66,738
$e^{SOC1}(157)$	93,469	5,121	23,068	2,169	2,142	116,536	120,848	67,048
$e^{SOC4}(183)$	94,363	5,066	22,818	2,170	2,141	117,181	121,492	66,478
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
$e^{Pre2008}(137)$	92,296	5,152	23,207	2,174	2,147	115,503	119,824	67,387
$e^{Post2008}(144)$	92,768	5,142	23,162	2,172	2,145	115,931	120,247	67,276
$e^{NR}(130)$	92,805	4,493	20,241	3,641	2,197	113,045	118,883	66,795
$e^{R}(130)$	92,805	4,493	20,241	3,641	2,197	113,045	118,883	66,795
$e^{VI}(130)$	92,805	4,493	20,241	3,641	2,197	113,045	118,883	66,795
$e^{IND}(130)$	92,805	4,493	20,241	3,641	2,197	113,045	118,883	66,795
$e^{SOC}(168)$	94,876	4,424	19,926	3,630	2,187	114,802	120,619	66,111
(M,M) Assortment: Three Musketeers and Milkyway								
$e^{SOC1}(154)$	94,316	4,454	20,063	3,632	2,189	114,379	120,200	66,398
$e^{SOC4}(178)$	95,161	4,398	19,812	3,631	2,186	114,972	120,789	65,878
$e^{Pre2008}(137)$	93,339	4,483	20,194	3,637	2,194	113,533	119,364	66,688
$e^{Post2008}(144)$	93,791	4,472	20,144	3,635	2,192	113,934	119,761	66,574
$e^{NR}(130)$	91,126	5,718	25,755	0	2,139	116,881	119,020	67,217
$e^{R}(130)$	91,126	5,718	25,755	0	2,139	116,881	119,020	67,217
$e^{VI}(130)$	91,126	5,718	25,755	0	2,139	116,881	119,020	67,217
$e^{IND}(130)$	91,126	5,718	25,755	0	2,139	116,881	119,020	67,217
$e^{SOC}(172)$	93,397	5,647	25,438	0	2,132	118,835	120,967	66,448
$e^{SOC1}(157)$	92,825	5,677	25,572	0	2,133	118,397	120,530	66,765
$e^{SOC4}(183)$	93,686	5,621	25,322	0	2,132	119,008	121,141	66,182
$e^{Pre2008}(137)$	91,673	5,708	25,713	0	2,137	117,387	119,523	67,111
$e^{Post2008}(144)$	92,139	5,698	25,668	0	2,135	117,807	119,942	66,998

Table A12: Simulated Profits with Weight on Consumer Surplus

Notes: Profit numbers represent the long-run expected profit from a top quartile machine with $MC = 0.15$ and $FC = 10$ but with weight of $\gamma = 3$ on consumer surplus ($\epsilon = -1$) in retailer's objective function. Retail profits do not include rebate payments. The socially-optimal assortment is (H, M). First column reports policy type and value in parenthesis. $FC = 10, MC = 0.15$.

References

BERNHEIM, B. D. AND M. WHINSTON (1998): "Exclusive Dealing," *Journal of Political Economy*, 106, 64–103.

BORK, R. (1978): *The Antitrust Paradox*, New York: Free Press.

CHAO, Y., G. TAN, AND A. C. L. WONG (2018): "All-units discounts as a partial foreclosure device," *The RAND Journal of Economics*, 49, 155–180.

KOLAY, S., G. SHAFFER, AND J. A. ORDOVER (2004): "All-Units Discounts in Retail Contracts," *Journal of Economics and Managment Strategy*, 13, 429–459.

O'BRIEN, D. P. (2013): "All-Units Discounts and Double Moral Hazard," FTC Working Paper No. 316.

POSNER, R. (1976): *Antitrust Law: An Economic Perspective*, University of Chicago Press.

WHINSTON, M. (1990): "Tying, Foreclosure, and Exclusion," *American Economic Review*, 80, 837–59.