

Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence

Christopher T. Conlon^y
Julie Holland Mortimer^z

June 7, 2017

Abstract

Vertical rebates are prominently used across a wide range of industries. These contracts may induce greater retail effort, but may also prompt retailers to drop competing products. We study these offsetting efficiency and foreclosure effects empirically, using data from one retailer. Using a field experiment, we show how the rebate allocates the cost of effort between manufacturer and retailer. We estimate models of consumer choice and retailer behavior to quantify the rebate's effect on assortment and retailer effort. We find that the rebate increases industry profitability and consumer utility, but fails to maximize social surplus and leads to upstream foreclosure.

We thank Mark Stein, Bill Hannon, and the drivers at Mark Vend Company for implementing the experiment used in this paper, providing data, and generally educating us about the vending industry. We thank Susanto Basu, Michaela Draganska, Scott Hemphill, Wills Hickman, J. F. Houde, Sylvia Hristakeva, Robin Lee, Joe Quinn, Michael Riordan, Greg Shafer, Mike Whinston, and seminar participants at Bates White, Boston College, Boston University, Harvard Business School, Harvard University (Economics Department), MIT, Pennsylvania State University, Rochester University, SUNY Stony Brook, Texas A & M University, Toulouse School of Economics, Tufts University, Universitat Autònoma de Barcelona, Universitat Pompeu Fabra, University of Toronto, University of Virginia, Washington University in St. Louis, Yale University, and participants at the Quantitative Marketing and Economics conference and the Workshop on Conditional Pricing Practices (held jointly by the Department of Justice and the Federal Trade Commission), for helpful comments. Any remaining errors are our own.

^yDepartment of Economics, NYU Stern, 44 W 4th St, New York City, NY 10012. email: cconlon@stern.nyu.edu

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

1 Introduction

Vertical arrangements between manufacturers and retailers have important implications for how markets function. These arrangements may align retailers' incentives with those of manufacturers, and induce retailers to provide demand-enhancing effort. However, they may also reduce competition, exclude competitors, and limit product choice for consumers. Many types of vertical arrangements can induce these offsetting efficiency and foreclosure effects, including resale price maintenance, exclusive dealing, vertical bundling, and rebates, among other contractual forms. Accordingly, these arrangements are a primary focus of antitrust authorities in many countries. Vertical rebates in particular are prominently used across a wide range of industries, including pharmaceuticals, hospital services, microprocessors, snack foods, and heavy industry, and have been the focus of several recent Supreme Court cases and antitrust settlements.¹

Although vertical rebate contracts are important in the economy and have the potential to induce both pro- and anti-competitive effects, understanding their economic impacts can be challenging. Tension between the potential for efficiency gains on one hand, and exclusion of upstream rivals on the other hand, implies that the contracts must be studied empirically in order to gain insight into the relative importance of the two effects. Unfortunately, the existence and terms of these contracts are usually considered to be proprietary information by their participating firms, frustrating most efforts to study them empirically. An additional challenge for analyzing the effect of vertical contracts is the difficulty in measuring downstream effort, both for the upstream firm and the researcher.

We address these challenges by examining a vertical rebate known as an All-Units Discount (AUD). The specific AUD we study is used by the dominant chocolate candy manufacturer in the United States: Mars, Inc.² The AUD implemented by Mars consists of three main features: a retailer-specific per-unit discount, a retailer-specific quantity target or threshold, and a 'facing' requirement that the retailer carry at least six Mars products.

¹Different forms of vertical rebates include volume-based discounts and 'loyalty contracts.' Volume-based discounts tie payments to a retailer's total purchases from the rebating manufacturer, but do not reference the sales of competing manufacturers. An all-units discount is a particular type of volume-based discount in which the discount is activated once sales exceed a volume threshold. Once activated, the discount applies retroactively to all units sold. We use the term 'loyalty contracts' to refer to payments that are calculated based on a retailer's sales volumes of both the rebating, and competing, manufacturers. Genchev and Mortimer (forthcoming) provides a review of empirical evidence on this class of contracts, including many of the relevant court cases.

²With revenues in excess of \$50 billion, Mars is the third-largest privately-held company in the United States (after Cargill and Koch Industries).

Mars' AUD stipulates that if a retailer meets the facing requirement and his total purchases exceed the quantity target, Mars pays the retailer an amount that is equal to the per-unit discount multiplied by the retailer's total quantity purchased. We examine the effect of the rebate contract through the lens of a retail vending operator, Mark Vend Company, for whom we are able to collect extremely detailed information on sales, wholesale costs, and contractual terms. The retailer also agreed to run a large-scale field experiment on our behalf, in which we exogenously remove two of Mars' best-selling products and observe subsequent substitution patterns, as well as the profit/revenue impacts for the retailer and all manufacturers. This provides important insight into the effect of the retailer's actions on manufacturer revenues, as well as the potential impact of the AUD on the retailer's decisions. To the best of our knowledge, no previous study has had the benefit of examining a vertical rebate contract using such rich data and exogenous variation.

The insights that we gain from studying Mars' rebate contract allow us to contribute to understanding principal-agent models in which downstream moral hazard plays an important role. Downstream moral hazard arises whenever a downstream firm takes a costly action that is beneficial to the upstream firm but not fully contractible. It is an important feature of many vertically-separated markets, and is thought to drive a variety of vertical arrangements such as franchising and resale price maintenance (RPM).³ However, empirically measuring the effects of downstream moral hazard is difficult. Downstream effort may be impossible to measure directly, and vertical arrangements are endogenously determined, making it difficult to identify the effects of downstream moral hazard on upstream firms. Our ability to exogenously vary the result of downstream effort (in this case, retail product availability), combined with detailed data on wholesale prices, allows us to directly document the effects of downstream moral hazard on the revenues of upstream firms.

In order to analyze the effect of Mars' AUD contract, we specify a model of consumer

of vertical restraints goes back at least to Telser (1960) and the *Downstream Moral Hazard* problem discussed in Chapter 4 of Tirole (1988).⁸ An important theoretical development on the potential foreclosure effects of vertical contracts is the so-called *Chicago Critique* of Bork (1978) and Posner (1976), which makes the point that because the downstream firm must be compensated for any exclusive arrangement, one should only observe exclusion in cases for which it maximizes the profits of the entire industry. Subsequent theoretical literature

on downstream moral hazard or effort decisions.¹⁵ The most closely-related empirical work is work on vertical bundling in the movie industry, and on vertical integration in the cable television industry. The case of vertical bundling, known as full-line forcing, is studied by Ho, Ho, and Mortimer (2012a) and Ho, Ho, and Mortimer (2012b), which examine the decisions of upstream firms to offer bundles to downstream retailers, the decisions of retailers to accept these 'full-line forces,' and the welfare effects induced by the accepted contracts. The case of vertical integration is studied by Crawford, Lee, Whinston, and Yurukoglu (2015), which examines efficiency and

and when substitute products or alternative distributors are not widely available." While the wide variety of arrangements and the diversity of market structures makes generalization difficult with any observed CPP (including the one we study here), the potential for both anti-competitive and efficiency effects makes it important to build on the empirical body of knowledge about these arrangements. As Genchev and Mortimer (forthcoming) point out, it is especially important to empirically analyze the impacts of CPPs that have not been selected through a process of litigation, to avoid selection bias in the set of contracts examined in the literature.

The rest of the paper proceeds as follows. Section 2 provides the theoretical framework for the model of retail behavior. Section 3 describes the vending industry, data, and the design and results of the field experiment, and section 4 provides the details for the empirical implementation of the model. Section 5 provides results, and section 6 concludes.

2 Theoretical Framework

2.1 Foreclosure and Optimal Assortments: A Motivating Example

We begin by providing a working definition, as well as some examples of the measures of *foreclosure* and *optimal assortment* to be used throughout the rest of our paper. To begin, we focus exclusively on the assortment decision (ignoring effort provision) of the downstream retailer (R) in response to a contract offered by a dominant upstream firm (M)

$(H;H)$. The dominant firm M offers the retailer R a transfer T in exchange for switching from $(H;H)$ to $(M;M)$. In order to make the retailer's decision non-trivial, we assume that $R(M;M) < R(H;H)$ (i.e., the retailer earns higher profits when stocking the rival's products).¹⁹ The following conditions (A1)-(A3) ensure that such a transfer is sufficient for M to foreclose its rival H .

$$(A1) \quad R + T > 0$$

$$(A2) \quad M - T > 0$$

$$(A3) \quad -H < M + R$$

(A1) specifies that the retailer prefers to switch from $(H;H)$ to $(M;M)$ after receiving a transfer of size T ; (A2), that the dominant firm would be willing to pay T to induce the retailers to switch from $(H;H)$ to $(M;M)$. The third assumption (A3) says that the profits lost by the rival H are smaller than those gained by M and R combined. Thus, (A3) guarantees that even if H offered its own transfer equal to its entire lost profits $-H(H;H)$, it could not prevent foreclosure.²⁰

The ability to obtain foreclosure as an equilibrium outcome is guaranteed by (A3), which may also be restated as $-H < R + M$. H is willing to give up all of her profits in order to avoid foreclosure. Thus, when foreclosure is observed, it must be the case that H 's losses are smaller than the gains of R and M combined. From the perspective of industry profits, $-H > 0$, we call this type of foreclosure 'industry optimal'.²¹

Adding a Third Assortment

Now we introduce a new assortment $(H;M)$ which yields intermediate profits for all players:

$$\begin{aligned} R(H;H) &> R(H;M) > R(M;M) \\ H(H;H) &> H(H;M) > H(M;M) \\ M(H;H) &< M(H;M) < M(M;M) \end{aligned} \tag{1}$$

¹⁹Under an AUD, the transfer would be conditional on meeting a quantity threshold or a facing requirement that is only satisfied under an $(M;M)$ assortment.

²⁰If H is fully excluded from the retailer shelf then $H(M;M) = 0$ and $-H = H(H;H)$.

²¹The effect of the change in assortment on consumer surplus $\Delta C > 0$ or overall social surplus, $\Delta C + \Delta I$ may differ from its effect for the industry.

For this case, we ignore the possibility of $(M;M)$, and introduce a new operator $H = (H;M) \quad (H;H)$, with the same set of assumptions:

$$(B1) \quad H^R + T_h = 0$$

$$(B2) \quad H^M - T_h = 0$$

$$(B3) \quad -H^H - H^M$$

Proofs in Appendix A.1.

The main takeaway is that M can set the vector of transfer payments T ; T_h ; and T_m in order to obtain full ($M;M$) or partial ($H;M$) foreclosure. We show that under (A1)-(A3), full foreclosure is feasible.²² However, if (B1)-(B3) and (C1), (C2), and (C4) 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 vector of transfers T , T_h and T_m , 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.

2.2 All Units Discount Rebates

In an All Units Discount (AUD) rebate, the transfer T to the retailer is calculated on the basis of all units sold to the retailer, conditional on obtaining a level of sales at or above a required threshold.

Assuming that the dominant manufacturer offers the same wholesale price across all

tion of M 's profits, rather than quantity purchased. Specifically,

$$\begin{aligned} & \delta \\ & < R(a) + d q_m(a) & \text{if } q_m(a) \geq \bar{q}_m \\ & : R(a) & \text{if } q_m(a) < \bar{q}_m \end{aligned}$$

cost of providing effort $c(e)$ is increasing in e . If we hold assortment fixed, the retailer's payoffs under the AUD, as a function of effort, are:

$$\begin{cases} & \delta \\ < & R(e) - c(e) + \beta^M(e) & \text{if } M(e) \geq \bar{M} \\ : & R(e) - c(e) & \text{if } M(e) < \bar{M} \end{cases} \quad (2)$$

The upstream firm can induce greater retailer effort via both features of the contract: (1) a larger per unit discount increases β so that R gives greater consideration to the profits of M ; (2) a larger choice of \bar{M} leads to greater retailer effort because $M(e)$ is increasing in effort. In our empirical example, we quantify both of these channels.

We provide a detailed solution to the effort problem in Appendix A.3. To summarize, when effort is non-contractible, R chooses one of three solutions to equation (2): either the interior solution to the effort problem with the rebate (the first line), which we denote e^R , the interior solution to the effort problem absent the rebate (the second line), which we denote e^{NR} , or the solution that makes the constraint bind, $\bar{e} : M(\bar{e}) = \bar{M}$. Thus, for $\bar{e} \leq e^R$, M can set the effort level of the retailer via the threshold \bar{M} , subject to satisfying the retailer's IR constraint. The set of effort levels that the threshold can target potentially includes the vertically-integrated, and the socially-optimal effort levels. Later, we characterize the critical values of \bar{M} in our empirical exercise.

An important consideration is whether the potential efficiency gains from increased retailer effort can offset the potential surplus lost from foreclosure. In order to analyze this question, we focus primarily on effort levels that maximize efficiency gains. One can examine the effort choice that is optimal for the bilateral/vertically-integrated firm $M + R$, which we denote e^{VI} , or for the industry (i.e., including profits of the rival), which we denote e^{IND} , or the effort level that maximizes social surplus, denoted e^{SOC} .

We enumerate these possibilities below:

$$\begin{aligned} e^{NR} &= \arg \max_e R(e) - c(e) \\ e^R &= \arg \max_e R(e) - c(e) + \beta^M(e) \\ e^{VI} &= \arg \max_e R(e) - c(e) + \beta^M(e), \end{aligned} \quad (3)$$

to the effort problem absent the rebate (the second line), -27(cial)-327(surplus,)-326(denote2)]T07F18 11.9552

$M(e)$ is increasing everywhere. This can be accomplished by choosing a threshold $\overline{M} > M(e^{VI})$. For $e < e^{VI}$ the bilateral surplus is increasing in e sort, and for $e > e^{VI}$ the bilateral surplus is decreasing in e sort; however, at all levels of e , e sort (weakly) functions as a transfer from R to M . Thus, in equilibrium, it may be possible to design a transfer that results in socially inefficient excess e sort.

3 The Vending Industry and Experimental Data

3.1 Data Description and Product Assortment

We observe data on the quantity and price of all products sold by one retailer, Mark Vend Company. Mark Vend is located in a northern suburb of Chicago, and services 728 snack machines throughout the greater Chicago metropolitan area.²⁷ Data are recorded internally at each of Mark Vend's machines, and include total vends and revenues since the last service visit to the machine. Any given snack machine can carry roughly 31 standard products at one time. These include salty snacks, cookies, and other products, in addition to 6-8 confection products.²⁸ We observe retail and wholesale prices for each product at each service visit during a 38-month panel for all snack machines in Mark Vend's enterprise. The dataset covers the period from January, 2006 through February, 2009. There is relatively little price variation over time for any given machine, and almost no price variation within a product category (e.g., confections) for a machine.

A focus in our empirical exercise is the set of products the retailer stocks in the last two slots in the confections category. Mark Vend chooses between stocking two additional Mars products (Milkyway and 3 Musketeers) or two Hershey Products (Reese's Peanut Butter Cups and Payday), or one product from each manufacturer. In table 1 we report the national

(around 52% of all confections sales). The non-Mars product most frequently stocked by Mark Vend is Nestle's Raisinets (at 47% of machine-weeks), which does not rank in the top 45 products nationally in confections sales.

There are two possible explanations for Mark Vend's departures from the national best-sellers. One is that Mark Vend has better information on the tastes of its specific consumers, and its product mix is geared towards those tastes. The alternative explanation is that the rebate induces Mark Vend to substitute from Nestle/Hershey brands to Mars brands when making stocking decisions, and that when Mark Vend does stock products from competing manufacturers (e.g., Nestle Raisinets), he chooses products that do not steal business from key Mars products.

3.2 Mars' AUD with Mark Vend

Mars' AUD rebate program is the most commonly-used vertical arrangement in the vending industry.²⁹ Under the program, Mars refunds a portion of a vending operator's wholesale cost at the end of a fiscal quarter if the vending operator meets a quarterly sales goal. The sales goal for an operator is typically set on the basis of its combined sales of Mars' products, rather than for individual Mars products. Mars' rebate contract also stipulates a minimum number of product 'facings' that must be present in an operator's machines, although in practice, this provision is difficult to enforce because Mars cannot observe the assortments in individual vending machines. The amount of the rebate and the precise threshold of the sales goal are specific to an individual vending operator, and these terms are closely guarded by participants in the industry.

We include some promotional materials from Mars' rebate program in figure 1.³⁰ The program employs the slogan *The Only Candy You Need to Stock in Your Machine!*, and specifies a facing requirement of six products and a quarterly sales target. The second page of the document shown in figure 1 refers to discontinuing a growth requirement, which we

²⁹For confections products, Mars is the dominant manufacturer in vending, and is the only manufacturer to offer a true AUD contract. The AUD is the only program offered to vendors by Mars. Hershey and Nestle offer wholesale 'discounts,' but these have a quantity threshold of zero (i.e., their wholesale pricing is equivalent to linear pricing). The salty snack category is dominated by Frito-Lay (a division of PepsiCo) which does not offer a rebate contract. We do not examine beverage sales, because many beverage machines at the locations we observe are serviced directly by Coke or Pepsi.

³⁰A full slide deck, titled '2010 Vend Program,' and dated December 21, 2009, is available at <http://vistar.com/KansasCity/Documents/Mars%202010%20Operatopr%20rebate%20program.pdf>. (Last accessed on April 19, 2015; available from the authors upon request.) These promotional materials represent the same type of rebate in which Mark Vend participated, but may differ from the terms available to Mark Vend during the period we study.

believe to be 5% (i.e., a target of 105% of year-over-year sales). On another page, not shown

from around 6.6 to around 5.3. Over the same time period, the number of Hershey facings increased from around 1 facing per machine to around 2 facings per machine. The right-hand-side panel of the table shows that the major switch was to swap Mars' Three Musketeers (stocked in around half of machines at the beginning of the sample) for Hershey's Reese's Peanut Butter Cups and Payday (stocked in 62% and 23% of machines respectively at the end of the sample period). Although it is difficult to attribute causality, it is worth pointing out that prior to the reduction in the threshold, both Reese's Peanut Butter Cups and Payday are effectively foreclosed, as they are stocked in very few of Mark Vend's machines.

product removals are recorded during each service visit.³⁴ Implementation of each product removal was fairly straightforward; the driver removed either one or both of the two top-selling Mars products from all machines for a period of roughly 2.5 to 3 weeks. The focal products were Snickers and Peanut M&Ms.³⁵ The dates of the product removal interventions range from June 2007 to September 2008, with all removals run during the months of May - October. Over all sites and months, we observe 185 unique products. We consolidate products that had very low levels of sales with similar products within a category that are produced by the same manufacturer, until we are left with the 73 'products' that form the basis of the rest of our exercise.³⁶

During each 2-3 week product removal period, most machines receive about three service visits. However, the length of service visits varies across machines, with some machines visited more frequently than others. Machines are serviced on different schedules, and as a result, it is convenient to organize observations by machine-week, rather than by visit, when analyzing the results of the experiment. When we do this, we assume that sales are distributed uniformly among the business days in a service interval, and assign those business days to weeks. Different experimental treatments start on different days of the week, and we allow our definition of when weeks start and end to depend on the client site

³⁴The machines are located in office buildings, and have substitution patterns that are very stable over time. In addition to the three treatments described here, we also ran several other treatment arms, for salty-snack and cookie products, which are described in Conlon and Mortimer (2010) and Conlon and Mortimer (2013b). The reader may refer to our other papers for more details.

³⁵Whenever a product was experimentally stocked-out, poster-card announcements were placed at the front of the empty product column. The announcements read "This product is temporarily unavailable. We apologize for any inconvenience." The purpose of the card was two-fold: first, we wanted to avoid dynamic effects on sales as much as possible, and second, Mark Vend wanted to minimize the number of phone calls received in response to the stock-out events. 'Natural,' or non-experimental, stock-outs are extremely rare for our set of machines. This implies that much of the variation in product assortment comes either from product rotations, or our own exogenous product removals. Product rotations primarily affect 'marginal' products, so in the absence of exogenous variation in availability, the substitution patterns between marginal products is often much better identified than substitution patterns between continually-stocked best-selling products. Conlon and Mortimer (2010) provides evidence on the role of the experimental variation for identification of substitution patterns.

³⁶For example, we combine Milky Way Midnight with Milky Way, and Ruffles Original with Ruffles Sour Cream and Cheddar. In addition to the data from Mark Vend, we also collect data on product characteristics online and through industry trade sources. For each product, we note its manufacturer,

and experiment.³⁷

Two features of consumer choice are important for determining the welfare implications of the AUD contract. These are, first, the degree to which Mark Vend's consumers prefer the marginal Mars products (Milky Way, Three Musketeers, Plain M&Ms) to the marginal Hershey products (Reese's Peanut Butter Cup, Payday), and second, the degree to which any of these products compete with the dominant Mars products (Peanut M&Ms, Snickers, and Twix). Our experiment mimics the impact of a reduction in retailer effort (i.e., restocking frequency) by simulating the stock-out of the best-selling Mars confections products. This provides direct evidence about which products are close substitutes, and how the costs of stock-outs are distributed throughout the supply chain. It also provides exogenous variation in the choice sets of consumers, which helps to identify the discrete-choice model of consumer choice.

In principle, calculating the effect of product removals is straightforward. In practice, however, there are two challenges in implementing the removals and interpreting the data generated by them. First, there is variation in overall sales at the weekly level, independent of our exogenous removals. Second, 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, and we rely on observational data for the control weeks. To mitigate these issues, we report treatment effects of the product removals after selecting control weeks to address these issues. We provide the details of this procedure in Appendix A.4.

3.4 Results of Product Removals

Our first exogenous product removal eliminated Snickers from all 66 vending machines involved in the experiment; the second removal eliminated Peanut M&Ms, and the third eliminated both products.³⁸ These products correspond to the top two sellers in the confections category, both at Mark Vend and nationwide.

One of the results of the product removal is that many consumers purchase another product in the vending machine. While many of the alternative brands are owned by Mars, several of them are not. If those other brands have similar (or higher) margins for Mark Vend, substitution may cause the cost of each product removal to be distributed unevenly across the supply chain. Table 5 summarizes the impact of the product removals for Mark

³⁷For example, at some site-experiment pairs, we define weeks as Tuesday to Monday, while for others we use Thursday to Wednesday.

³⁸As noted in table 1, both Snickers and Peanut M&Ms are owned by Mars.

Vend. In the absence of any rebate payments, we see the following results. Total vends decrease by 217 units and retailer profits decline by \$56.75 when Snickers is removed. When Peanut M&Ms is removed, vends go down by 198 units, but Mark Vend's average margin on all items sold in the machine rises by 0.78 cents, and retailer revenue declines only by \$10.74 (a statistically insignificant decline). Similarly, in the joint product removal, overall vends decline by roughly 283 units, but Mark Vend's average margin rises by 1.67 cents per unit, so that revenue declines by only \$4.54 (again statistically insignificant).³⁹

Table 6 examines the impact of the product removals on the upstream firms. Removing Peanut M&Ms decreases Mars' revenue by about \$68.38, compared to Mark Vend's loss of \$10.74; thus roughly 86.4% of the cost of stocking out is born by Mars (reported in the fifth column). In the double removal, because Peanut M&M customers can no longer buy Snickers, and Snickers customers can no longer buy Peanut M&Ms, Mars bears 96.7% of the cost of the stockout. In the Snickers removal, most of the cost appears to be born by the downstream firm; one potential explanation is that among consumers who choose another product, many select another Mars Product (Twix or Peanut M&Ms). We also see the impact of each product removal on the revenues of other manufacturers. Hershey (which owns Reese's Peanut Butter Cups and Hershey's Chocolate Bars) enjoys relatively little substitution in the Snickers removal, in part because Reese's Peanut Butter cups are not available as a substitute. In the double removal, when Peanut Butter Cups are available, Hershey profits rise by nearly \$61.43, capturing about half of Mars' losses. We see substitution to the two Nestle products in the Snickers removal, so that Nestle gains \$19.32 as consumers substitute to Butterfinger and Raisinets; Nestle's gains are a smaller percentage of Mars' losses in the other two removals.

Direct analysis of the product removals can only account for the marginal cost aspect of the rebate (i.e., the price reduction given by \bar{p}); one requires a model of restocking in order to account for the threshold aspect, \bar{M} . By more evenly allocating the costs of stocking out, the rebate should better align the incentives of the upstream and downstream firms, and lead the retailer to increase his overall service level. Returning to table 5, the right-hand panel reports the retailer's profit loss from the product removals after accounting for his rebate payments, assuming he qualifies. We see that the rebate reallocates approximately (\$17, \$30, \$50) of the cost of the Snickers, Peanut M&Ms, and joint product removals from the upstream to the downstream firm. The last column of table 6 shows that after accounting

³⁹Total losses appear smaller in the double-product removal in part because we sum over a smaller sample size of viable machine-treatment weeks (89) for this experiment, compared to the Peanut M&Ms removal (with 115 machine-treatment weeks).

for the rebate contract, the manufacturer bears about 50% of the cost of the Peanut M&Ms removal, 60% of the cost of the joint removal, and 12% of the cost of the Snickers removal.

4 Estimation

4.1 Consumer Choice

In order to consider the optimal product assortment, we need a parametric model of consumer choice that predicts sales for a variety of different product assortments. We estimate a mixed (random-coefficients) logit model on our sample of 66 machines (including both experimental and non-experimental periods).⁴⁰

We consider a model of utility in which consumer i receives utility from choosing product j in market t of:

$$U_{ijt} = \beta_j + \gamma_{ijt} + \epsilon_{ijt} \quad (4)$$

The parameter β_j is a product-specific intercept that captures the mean utility of product j in market t , and γ_{ijt} captures individual-specific correlation in tastes for products. Each consumer has an outside option $U_{i0t} = \epsilon_{i0t}$

t . The first allows for 15,256 fixed effects, at the level of a machine-service visit, while the second allows for 2,710 fixed effects, at the level of a machine-choice set (i.e., we combine machine-service visit 'markets' for which the choice set does not change). We report the log-likelihood, Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) for each specification. We use BIC to select the specification with 2,710 t fixed effects. Our simulated MLE parameters tend to be very precisely estimated, because we observe 2.96 million sales.

Parametric identification of d_j and β parameters is straightforward. The d_j parameters would be identified from average sales levels in even a single market after we normalize the utility of the outside good to zero. (s38 tnm(Info4055.imeormat92.96)]TJ 03lma.928 Td [(mil7.633.9

The retailer's problem is:

$$\max_{a:e} \begin{cases} R(a:e) - c(e) + M(a:e) & \text{if } M(a:e) \geq \bar{M} \\ R(a:e) - c(e) & \text{if } M(a:e) < \bar{M} \end{cases} \quad (7)$$

where $R(a:e)$ is the variable profit of the retailer absent any rebate payment, $M(a:e)$ is the variable profit of the dominant manufacturer M , and $c(e)$ represents the cost of retailer effort.

The retailer's assortment decision involves simple discrete comparisons across a finite number of choices. We explain the set of potential assortments that we analyze in section 4.2.3. For each potential choice of assortment, we calculate the retailer's optimal choice of effort.

4.2.1 Retail Effort Choice: Dynamic Model of Re-stocking

We believe that MarkVend's effort decision is operationalized as follows. At the beginning of each quarter, MarkVend decides on an (enterprise-wide) policy to restock after e likely consumers have arrived at all of his vending machines.⁴⁶ He then translates this policy into a restocking schedule for each individual vending machine (e.g., every Tuesday, every 10 days, every other day, etc.) based on knowledge of a machine-specific arrival rate. Once the schedule for the quarter is set, he breaks up the schedule into individual service routes, and assigns routes to drivers and trucks. In order to reduce the number of consumer arrivals between service visits, MarkVend must hire additional trucks and drivers, which increases his costs. An implication of this setup is that MarkVend commits to a restocking policy for an entire quarter. This means that if sales are below expectations (i.e., if he repeatedly draw from the left-tail of the consumer arrival distribution), MarkVend does not adjust his stocking policy until the next quarter.⁴⁷

In our application, we consider the specific case in which the retailer chooses the restocking frequency. We model the retailer's choice of effort, e , using an approach similar to Rust

⁴⁶Mars' AUD rebate contract is evaluated quarterly on the basis of MarkVend's entire enterprise, which includes 728 snack vending machines.

⁴⁷Within a quarter, it appears as the most machines are on an extremely predictable fixed schedule, and there is no evidence that the schedule is adjusted in either direction towards the end of each quarter. This is consistent with a model of effort in which the frequency of service is set in response to the payoff function, but the schedule is not set dynamically within a quarter as a function of the distance from the threshold. As MarkVend does not observe sales, except at the time of a service visit, this makes a lot of sense. He doesn't have new information by which to dynamically adjust a service schedule across days.

(1987), but 'in reverse.' Rather than assuming that observed retailer wait times are optimal and using Rust's model to estimate the cost of re-stocking, we use an outside estimate of the cost of re-stocking based on wage data from the vending operator, and use the model to

This also enables us to evaluate profits under alternative stocking policies x^j , or policies that arise under counterfactual market structures. For example, in order to understand the incentives of a vertically-integrated firm, $M + R$, we can replace $u(x)$ with $(u^R(x) + u^M(x))$, which incorporates the profits of the dominant upstream manufacturer. Likewise, we can consider the industry-optimal policy by replacing $u(x)$ with $(u^R(x) + u^M(x) + u^H(x) + u^N(x))$.

To find the optimal policy we iterate between (9) and the policy improvement step:

$$x = \min_x : u(x) = FC + \sum_j P(x^j|x) V(x^j;x) \quad (10)$$

The fixed point $(x; V(x;x))$ maximizes the long-run average profit of the agent $\sum_j P(x^j|x) V(x^j;x)$ where $P = \sum_j P(x^j|x)$ is the ergodic distribution corresponding to the post-decision transition matrix. These long-run profits will become the basis on which we compare contracts and product assortment choices.

4.2.2 Retail Effort Choice: Empirical Implementation

In order to compute the dynamic restocking model, we construct a 'representative vending machine' via the following procedure. We define a 'full machine' as one that contains a set of the 29 most commonly-stocked products, which we report in table 8, and we use actual machine capacities for each product.⁵⁰ Beginning with a full machine, we simulate consumer arrivals one at a time and allow consumers to choose products in accordance with the mixed logit choice probabilities $s_{jt}(x; a_t)$ (including an outside option of no-purchase). After each consumer choice, we update the inventories of each product and adjust the set of available products a_t if a product has stocked out. When products stock out, consumers substitute to other products, including the no-purchase option. We continue to simulate consumer arrivals until the vending machine is empty. We average over 100,000 simulated chains to construct the expected profits after x consumers have arrived, and fit a smooth Chebyshev polynomial to the profits of each agent $u^R(x); u^M(x); u^H(x); u^C(x)$.⁵¹

The state variable of our dynamic programming problem, X_t , is the number of potential consumers who have arrived since our 'representative vending machine' was last restocked. The exogenous state transition matrix $P(X_{t+1} = X_t + j | X_t) = P(j | X_t)$ is the incremental number of potential consumers who arrive to the representative vending machine each business day.

⁵⁰These capacities are nearly uniform across machines, and are: 15-18 units for each confection product, 11-12 units for each salty snack product, and around 15 units for each cookie/other product.

⁵¹The fit of the 10th order Chebyshev polynomial is in excess of $R^2 = 0.99$. It is generally well behaved except at the very edges of the state space, but these are far from our optimal policies.

We assume that the arrival rate has a discrete distribution.⁵² In a separate stage, we use 28 of our 66 experimental machines to form a non-parametric estimate of $P(x)$. These 28 machines have an average daily sales volume of 15.1 units and a standard deviation of 2.0 units.⁵³ For each service-visit observation at each of these machines, we use the number of estimated consumer arrivals since the last service visit, and divide this by the number of elapsed business days since the last visit to compute the number of daily consumer arrivals, x_t .⁵⁴ E ort policies are not particularly sensitive to the specification of the arrival process.⁵⁵

We choose a daily discount factor $\beta = 0.999863$, which corresponds to a 5% annual interest rate. We assume a fixed cost of a restocking visit, $FC = \$10$, which approximates the per-machine restocking cost using the driver's wage and average number of machines serviced per day. As a robustness test, we also consider $FC = \text{€}5.15g$, which generate qualitatively similar predictions. In theory, one should be able to estimate FC directly from the data using the technique of Hotz and Miller (1993). However, our retailer sets a level of service that is too high to rationalize with any optimal stocking behavior, often restocking a day before any products have stocked-out.⁵⁶ This is helpful as an experimental control, but makes identifying FC from data impossible.

In order to speed up computation, we normalize our state space when solving the dynamic programming problem. Instead of working with the number of consumers to arrive at

⁵²This mimics Rust (1987) who estimates a discrete distribution of weekly incremental mileage.

⁵³The machines in this group have higher than average sales volumes, but are not the largest machines. We chose this group for our exercise because we think it is the most important set of machines for determining the retailer's re-stocking decision. For additional detail, please see Appendix A.5.

⁵⁴Note that the data report average daily sales, rather than consumer arrivals (i.e., there are no cameras on the vending machines). As in the consumer choice model, the relationship between observed sales and consumer arrivals depends on availability. If a machine is empty, no sales will occur, regardless of the consumer arrival rate. The consumer choice model adjusts for this by allowing substitution to remaining products (including the outside good) when a machine is not fully stocked. Our estimate of consumer arrivals uses the same adjustment.

⁵⁵Doubling or tripling the rate at which consumers arrive has very little effect on the optimal e ort policy, because policies are defined in terms of the cumulative number of consumer arrivals (rather than days, for example). In robustness tests, we assume that the firm can make decisions consumer-by-consumer, or only every four `days.' With appropriate scaling of the discount factor β , the optimal policies change by only 2-3 units.

⁵⁶In conversations with the retailer about his service schedule, he provided two explanations of this fact. First, he suspected that he was over-servicing, and reduced service levels after our field experiment. Second, he explained that high service levels are important to obtaining long-term (3-5 year) exclusive service contracts with locations. Our specific experimental locations almost certainly do not reflect a company-wide servicing policy. Specifically, these are high-end office buildings with high service expectations. Public

the vending machine, we work with the number of consumers who would have likely made a purchase at a hypothetical 'full' vending machine. This saves us from simulating large numbers of consumers who always choose the outside good, independent of product assortment. We thus label our state-space as 'likely' consumer arrivals instead of 'potential' consumer arrivals from this point forward.⁵⁷

By simulating from our consumer choice model in section 4.1, we can compute the payoffs to each agent from any assortment a and any effort level e using equation (9). For the retailer, with effort policy e :

$$R(a;e) = \frac{P(e)}{(1 - P(e))^{-1} \delta^R(x;a)}; \quad (11)$$

and represents the net present value of the long-run average (infinite horizon) profits of a single representative vending machine under assortment a

payoffs at (a,e) for each agent for 15 possible assortments. Each of the 15 possible assortments includes Mark Vend's five most commonly-stocked chocolate confections products: four Mars products (Snickers, Peanut M&Ms, Twix, and Plain M&Ms), and Nestle's Raisinets. The retailer is always worse off if he replaces any of these five products with a different product. We then allow the retailer to choose any pair of products for the final two slots in the confections category from a set of six products. The six products we consider include two Mars products (Milky Way and Three Musketeers), two Hershey products (Reese's Peanut Butter Cup and PayDay), and two Nestle products (Butterfinger and Crunch).⁵⁹ Although we compute the full model for all 15 possible assortments, only three end up being payoff-relevant: (M;M) { 3 Musketeers and MilkyWay, (H;M) { 3 Musketeers and Reese's Peanut Butter Cups, and (H;H) { Reese's Peanut Butter Cup and PayDay.

Finally, Mark Vend's assortment decision is discrete (either a product is on the shelf of the vending machine or it isn't), and our effort decision is discrete (we are restricted to restocking after an integer number of likely consumer arrivals). Thus, we can, and do,

order to convert consumer surplus into dollars, we perform a calibration exercise in which we assume that the median own-price elasticity is -2 . We view this as a relatively inelastic

rebate at the observed $\pi^M(H;M) = 1,882$ exceeds the gains to Mars ($\pi^M = 1,657$). Thus, Mars pays more to partially foreclose Hershey than it expects to gain from partial foreclosure. This cannot be an equilibrium outcome.

The second column of the second pane of table 9 starts from $(H;M)$ and considers a move to $(M;M)$. Now Reese's Peanut Butter Cup is replaced by MilkyWay and Hershey is fully foreclosed. Again, the retailer gives up some profit absent the rebate payment ($\pi^R = 308$), the dominant firm gains ($\pi^M = 1,338$), and bilateral surplus increases ($\pi^M + \pi^R = 1,030$). However, the gain in bilateral surplus is smaller than the losses to the rival (π^H

would be too generous. Likewise, if Mars believes that, absent the rebate, the retailer would have stocked $(H;M)$, the rebate would not be generous enough to induce the retailer to switch from $(H;M)$ to $(M;M)$.

5.2 Role of the Threshold

These results are meant to parallel those in section 2.3. We explore how the rebate threshold τ^M affects the retailer's choice of assortment and effort, assuming that wholesale prices and the rebate discount are fixed at their observed values. Figure 2 plots two curves. Each curve represents the profits of the retailer after receiving the rebate (i.e., $R(a;e) + \tau^M(a;e)$). The horizontal axis reports revenue of the dominant firm, R^M . The left curve represents the retailer's profits with an $(H;M)$ assortment. The right curve represents the retailer's profits with a $(M;M)$ assortment. As we move across each curve from left to right, the retailer's

until he reaches $^{-M} = 13$;

order levels.⁶⁷ Under an $(M;M)$ assortment, switching from the no-rebate retailer order policy (e^{NR}) to the vertically-integrated optimal order policy (e^{VI}) increases the restocking frequency by 7.95%. We use the e^{NR} order level rather than the e^R order level as our baseline in order to capture the maximum potential efficiency gains from the rebate contract.⁶⁸ Higher order is costly to the retailer ($\Delta M = -55$) and beneficial to Mars ($\Delta M = 128$), leading to a net gain in producer surplus ($\Delta PS = 63$) once we include competing manufacturers. Most

socially-optimal effort level, e^{SOC}

The right panel of table 14 conducts the same exercise, but assumes that the retailer would choose the $(H;M)$ assortment in the absence of Mars' rebate. Under this scenario, the rebate is much too generous and could be reduced by 38.18% to 44.79% while still foreclosing Hershey. Relatedly, holding θ fixed, Hershey would need to set a negative wholesale price (i.e., pay the retailer to sell its products). This highlights the fact that the rebate terms are only sensible as a device to make the retailer switch from $(H;H)$ to $(M;M)$.⁷⁰

5.5 Comparison to Uniform Wholesale Pricing by Mars

In lieu of an AUD, Mars could charge a lower wholesale price without conditioning on a threshold θ^M . Table 15 presents results for a uniform wholesale price by Mars. We hold θ fixed the wholesale prices of Hershey and Nestle $(w_h; w_n)$, and compute a new optimal wholesale price for M , w_m^U . The resulting set of wholesale prices $(w_m^U; w_h; w_n)$ does not constitute an equilibrium (because $(w_h; w_n)$ are not allowed to adjust). Therefore, the exercise is meant as tool to understand how the AUD reduces the price of foreclosure to the dominant firm, rather than reflecting what would happen to equilibrium prices in the absence of an AUD by Mars.

The main result is that Mars' wholesale price is lower than the post-discount wholesale price under the AUD. Effectively, Mars pays more for foreclosure without the threshold. We quantify exactly how much more by comparing Mars' uniform wholesale price to an AUD that forecloses under two different effort levels: e^R and e^{VI} . Mars' profit (after rebates), $(1 - \beta) \pi^M$, falls from \$11,005 to \$10,094, for a loss of \$911. The retailer's profit increases by a similar amount (\$921, or \$39,103 - \$38,182). The gains to the retailer are slightly larger if the threshold under the AUD had been used to implement the vertically-integrated effort level.

We plot the best response of Mars to the observed $(w_h; w_n)$ prices in figure 4. We do not consider an equilibrium in which all three upstream firms simultaneously set wholesale prices $(w_m; w_h; w_n)$. The challenge for modeling this is that no Nash equilibrium exists in pure strategies because the retailer's assortment decision is discrete; only a mixed-strategy Nash equilibrium exists.⁷¹

⁷⁰It should also be clear that adjusting the baseline from $(H;H)$ to $(H;M)$ means that the current rebate violates Mars's IR constraint (B2) as noted in table 9.

⁷¹The non-existence of pure-strategy equilibria is well documented in the theoretical literature (e.g., see recent work by Jeon and Menicucci (2012)), and derives from the fact that agents' best-response functions are discontinuous, and need not cross. The mixed-strategy Nash equilibrium of the uniform wholesale pricing

5.6 Implications for Mergers

Vending is one of many industries for which retail prices are often fixed across similar products and under different vertical arrangements. Indeed, there are many industries for which the primary strategic variable is not retail price, but rather a slotting fee or other transfer payment between vertically-separated firms. Thus, our ability to evaluate the impact of a potential upstream merger may turn on how the merger affects payments between firms in the vertical channel. We consider the impact of three potential mergers (Mars-Hershey, Mars-Nestle, and Hershey-Nestle) on the AUD terms offered to the retailer by Mars. Given the degree of concentration in the confections industry, antitrust authorities would likely investigate proposed mergers, especially mergers involving Mars.⁷²

Table 16 measures how competing manufacturers might respond to an upstream merger. The first column duplicates the second column of table 14 as a baseline. In the second column, we examine a potential Mars-Hershey merger. We assume that after the merger, the Hershey product (Reeses Peanut Butter Cup) is priced at the Mars wholesale price and included in Mars' rebate contract. The merged (Mars-Hershey) firm is now happy for consumers to substitute to Reese's Peanut Butter Cups, and the AUD is able to achieve the industry-optimal (and socially-optimal) product assortment of (H,M).⁷³ The merged firm faces competition from Nestle (Nestle Crunch and Butterfinger), which charges lower wholesale prices but sells less-popular products.⁷⁴ In the absence of an AUD, the retailer maximizes his profit by stocking the two Nestle products, but the AUD induces the retailer to choose an (H,M) assortment and an e^{VI} effort policy (evaluated at the observed discount and wholesale prices).⁷⁵ Although not reported in table 16, the AUD also maximizes social surplus by inducing the socially-optimal (H;M) assortment and a high effort level.⁷⁶

We consider the possibility that Nestle may be able to cut its price in order to avoid having Butterfinger and Nestle Crunch foreclosed. Following the same exercise that was performed in table 14, we find that Nestle would need to charge a negative wholesale price to the retailer in order to induce him to stock the less-popular Nestle products (similar to condition (A3)). Knowing that the Nestle products provide weak discipline for the merged Mars-Hershey firm, we next examine whether the merged Mars-Hershey firm can reduce the

⁷²For a related analysis of diversion ratios in this market, see Conlon and Mortimer (2013b).

⁷³We assume that the AUD retains \bar{a} at the pre-existing level, and sets $\bar{M} = M(e^{VI}(H;M))$ to induce the vertically-integrated optimal level of effort.

⁷⁴We use Nestle's observed wholesale price when computing changes in profits and producer surplus.

⁷⁵Table 16 reports changes in variable profit for each agent, but not levels. For the full details of post-merger profits (or revenues for manufacturers) at all $(a;e)$, please see Appendix A.9.

⁷⁶Producer surplus and consumer utility for each potential merger are also reported in Appendix A.9.

generosity of its rebate. Pre-merger, we found that the rebate was 3.53% too generous; after the merger it is 42.3% too generous. This implies that the market is unlikely to obtain the post-merger outcome in which the retailer stocks the socially-optimal assortment, because the discipline imposed by Nestle's products is likely too weak to keep the current AUD terms in force.

We perform a similar exercise in the third column, in which we allow Mars and Nestle to merge. The main difference now is that the merged firm internalizes the profits of Nestle's Raisinets, and is able to include the profit from Raisinets in the rebate. This again provides incentives for the merged firm to reduce the generosity of the rebate (by 12.67%).⁷⁷ Finally, we examine a Hershey-Nestle merger in the final column. Giving Hershey access to the profits of Raisinets does very little, because Raisinets is not in danger of being foreclosed. This exercise closely resembles our baseline (No Merger) scenario.

Throughout the paper, we report the variable profits for the retailer; it is likely that his overall operating profits, after accounting for administrative and overhead costs, are substan-

Identification of both the consumer choice and retailer-effort models benefits from exogenous

References

Aghion, P., and P. Bolton (1987): "Contracts as Barriers to Entry," *American Economic Review*, 77(3), 388{401.

Asker, J. (2016): "Diagnosing Foreclosure due to Exclusive Dealing," *Journal of Industrial Economics*, 64(3), 375{410.

Asker, J., and H. Bar-Isaac (2014): "Raising Retailers' Profits: On Vertical Practices and the Exclusion of Rivals," *American Economic Review*, 104(2), 672{686.

Baker, G. P., and T. N. Hubbard (2003): "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information," *American Economic Review*, 93(3), 551{572.

Bernheim, B. D., and M. Whinston (1998): "Exclusive Dealing," *Journal of Political Economy*, 106(1), 64{103.

Berry, S.

Whinston, M. (1990): "Timing, Fairness, and Exclusion,"
Shepard, A. (1993): "Contractual Form, Retail Price, and Asset Characteristics in Gasoline
Retailing," *RAND Journal of Economics*, 24(1), 58{77.

Sinkinson, M. (2014): "Pricing and Entry Incentives with Exclusive Contracts: Evidence
from Smartphones," Working Paper.

Stokey, N. L., R. Lucas, and E. Prescott (1989): *Recursive Methods for Economic
Dynamics*. Harvard University Press.

Sudhir, K., and V. R. Rao (2006): "Do Slotting Allowances Enhance Efficiency or Hinder
Competition?," *Journal of Marketing Research*, 43(2), 137{155.

Telser, L. (1960): "Why Should Manufacturers Want Fair Trade?," *from 222(Ass)-1(et)-222(Characte*

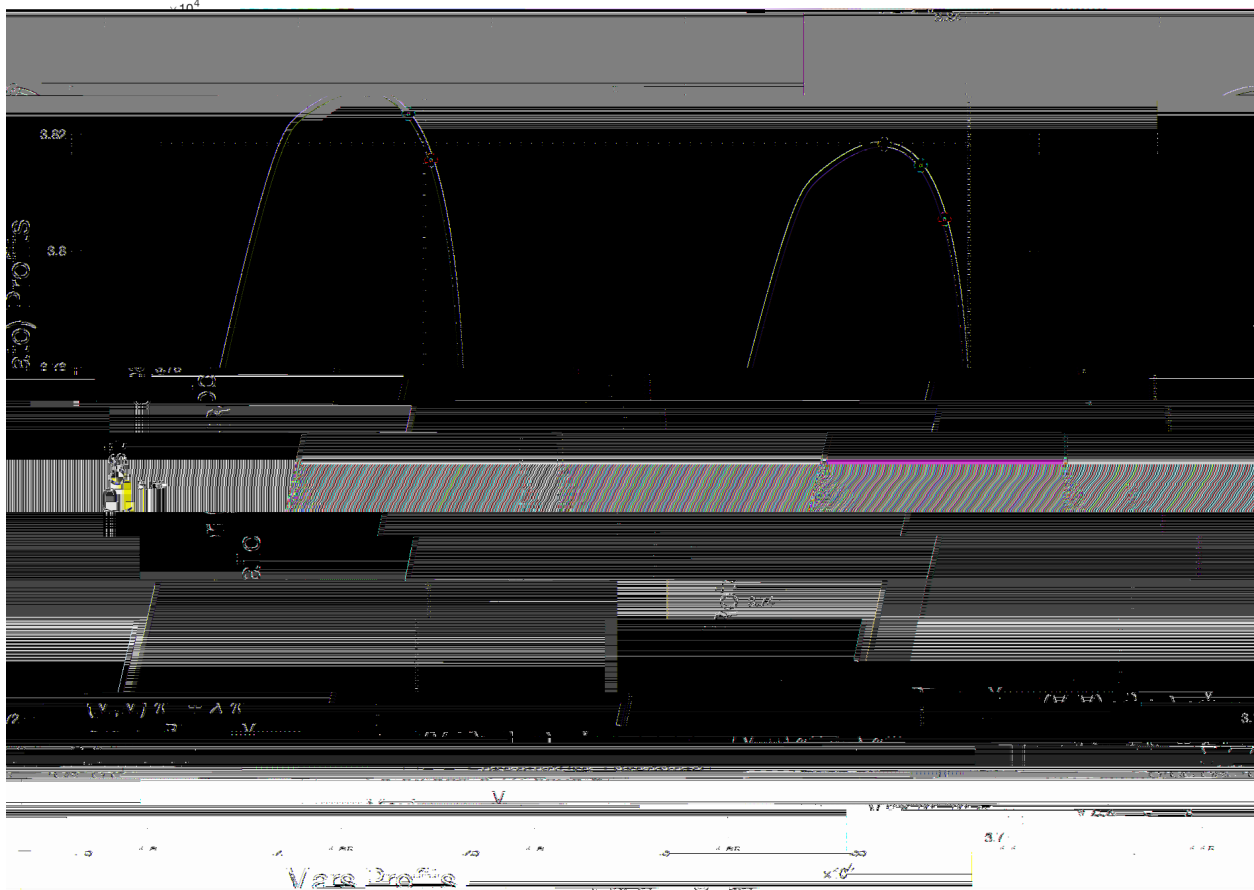
Train, K. (2003): *Discrete Choice Methods with Sim*

Figure 1: Mars Vend Operator Rebate Program




Notes: From '2010 Vend Program' materials, dated December 21, 2009; last accessed on February 2, 2015 at <http://vistar.com/KansasCity/Documents/Mars%202010%20Operatopr%20rebate%20program.pdf>.

Figure 2: Impact of AUD Quantity Threshold on Retail Assortment Choice



Notes: Figure reports retailer variable profit under two assortment choices ((H,M) on the left and (M,M) on the right), against revenues of Mars products. For a threshold $\bar{M} = 11.912$ (noted by the vertical dashed line), the retailer prefers to switch his assortment from (H,M) to (M,M). Three points are marked on each curve. The left-most point on each curve represents an e^R assortment policy for the relevant assortment; the point to the right of e^R represents an e^{VI} assortment policy, and the right-most point represents an e^{SOC} assortment policy.

Figure 3: Profits Per Consumer as a Function of the Restocking Policy



Notes: Each curve reports the profits of the retailer, Mars, Hershey and Nestle as a function of the retailer's restocking policy, using the product assortment in which the retailer stocks 3 Musketeers (Mars) and Reese's Peanut Butter Cups (Hershey) in the final two slots. Specifically, the vertical axes report variable profit per consumer for each of the four firms, and the horizontal axes report the number of expected sales between restocking visits.

Figure 4: Mars Profits as a Function of Price (Linear Pricing)

Notes: Reports Mars' profit at different linear wholesale prices, holding fixed the wholesale prices of Hershey and Nestle. The discontinuities reflect prices at which the retailer drops a Mars product from its assortment.

Table 1: Comparison of National Availability and Shares with Mark Vend

Manufacturer	Product	Rank	National:		Mark Vend:		Experimental:	
			Availability	Share	Availability	Share	Availability	Share
Mars	Snickers	1	89	12	87	16.9	97	21.3
Mars	Peanut M&Ms	2	88	10.7	89	16.0	97	22.1
Mars	Twix Bar	3	67	7.7	80	12.6	79	13.0
Hershey	Reeses Peanut Butter Cups	4	72	5.5	71	6.6	45	6.2
Mars	Three Musketeers	5	57	4.3	35	3.1	41	5.2
Mars	Plain M&Ms	6	65	4.2	71	6.6	45	6.2
Mars	Starburst	7	38	3.9	41	3.2	16	1.0
Mars	Skittles	8	43	3.9	65	5.6	79	6.3
Nestle	Butterfinger	9	52	3.2	32	2.1	32	2.6
Hershey	Hershey with Almond	10	39	3	1	0.1	0	0.0
Hershey	PayDay	11	47	2.9	13	1.2	1	0.1
Mars	Milky Way	13	39	1.7	33	2.8	18	1.5
Nestle	Raisinets	> 45	N/R	N/R	45	4.0	81	8.7

Notes: National Rank, Availability and Share refers to total US sales for the 12 weeks ending May 14, 2000, reported by Management Science Associates, Inc., at <http://www.allaboutvend.com/studies/study2.htm>, accessed on June 18, 2014. National figures are not reported for Raisinets because they are outside of the 45 top-ranked products. By manufacturer, the national shares of the top 45 products (from the same source) are: Mars 52.0%, and Hershey 20.5%. For Mark Vend, shares are: Mars 73.6%, and Hershey 15.0% and for our experimental sample Mars 78.3% and Hershey 13.1% (calculations by authors).

Table 2: Assortment Response to Changes in the Threshold

	Achieved Threshold %	Total Vends	Mars Share
2007q1	109.16	1000.00	20.20
2007q2	106.29	1087.45	19.77
2007q3	100.81	1008.57	20.94
2007q4	105.23	1092.49	19.97
2008q1	106.27	1103.42	19.45
2008q2	97.20	1057.32	19.77
2008q3	91.88	1014.13	19.14
2008q4	87.02	1048.26	18.11
2009q1	87.03	1058.54	17.65

Notes: Achieved threshold % reports the ratio of total Mars sales relative to Mars sales in the same quarter one year prior. For quarters 2007q1-2008q1 we believe the target to be 100% with a bonus payment at 105%. For quarters 2008q3-2009q1 we believe the threshold was reduced to 90%.

Table 3: Average Number of Confections Facings Per Machine-Visit

	Mars	Hershey	Nestle	Mars		Hershey	
				Milkyway	3 Musketeer	PB Cup	Payday
2006q1	6.64	1.32	2.05	0.26	0.50	0.19	0.08
2006q2	6.70	1.06	2.02	0.26	0.49	0.15	0.03
2006q3	6.76	0.81	2.02	0.29	0.56	0.03	0.01
2006q4	6.74	0.85	2.00	0.31	0.55	0.01	0.04
2007q1	6.61	1.13	1.58	0.32	0.56	0.00	0.08
2007q2	6.24	1.44	1.17	0.31	0.53	0.00	0.18
2007q3	6.21	1.63	1.08	0.29	0.54	0.01	0.21
2007q4	6.26	1.73	1.03	0.30	0.51	0.15	0.20
2008q1	5.98	2.08	0.97	0.38	0.29	0.51	0.19
2008q2	5.57	2.29	0.93	0.43	0.03	0.66	0.21
2008q3	5.37	2.29	0.91	0.41	0.00	0.63	0.23
2008q4	5.48	2.19	0.89	0.40	0.01	0.62	0.24

Table 7: Random Coefficients Choice Model

	Parameter Estimates	
<i>Salt</i>	0.506 [.006]	0.458 [.010]
<i>Sugar</i>	0.673 [.005]	0.645 [.012]
<i>Peanut</i>	1.263 [.037]	1.640 [.028]
# Fixed Effects t	15,256	2,710
LL	-4,372,750	-4,411,184
BIC	8,973,960	8,863,881
AIC	8,776,165	8,827,939

Notes: The random coefficients estimates correspond to the choice probabilities described in section 4, equation 5. Both specifications include 73 product fixed effects. Total sales are 2,960,315.

Table 8: Products Used in Counterfactual Analyses

`Typical Machine' Stocks:	
Confections:	Salty Snacks:
Peanut M&Ms	Rold Gold Pretzels
Plain M&Ms	Snyders Nibblers
Snickers	Ruffles Cheddar
Twix Caramel	Cheez-It Original
Raisinets	Frito
Cookie:	Dorito Nacho
Strawberry Pop-Tarts	Cheeto
Oat 'n Honey Granola Bar	Smartfood
Grandma's Chocolate Chip Cookie	Sun Chip
Chocolate Chip Famous Amos	Lays Potato Chips
Raspberry Knotts	Baked Lays
Other:	Munchos Potato Chips
Ritz Bits	Hot Stuff Jays
Ruger Vanilla Wafer	
Kar Sweet & Salty Mix	
Farley's Mixed Fruit Snacks	
Planter's Salted Peanuts	
Zoo Animal Cracker Austin	

Notes: These products form the base set of products for the `typical machine' used in the counterfactual exercises. For each counterfactual exercise, two additional products are added to the confections category, which vary with the product assortment selected for analysis.

Table 9: Assortment Decisions with Fixed E ort

	(H,H)	(H,M)	(M,M)
e^R	257	261	259
R	36,656	36,394	36,086
M	1,617	1,882	2,096
M	10,106	11,763	13,101
H	2,167	1,299	0
$R + M$	46,762	48,157	49,187
$R + M + H$	48,929	49,456	49,187
from	$(H:H)$	$(H:M)$	$(H:H)$
to	$(H:M)$	$(M:M)$	$(M:M)$
R	-262	-308	-570
	[2.50]	[0.95]	[2.44]
M	1,657	1,338	2,995
	[17.64]	[4.25]	[21.09]
$M+R$	1,395	1,030	2,425
	[15.63]	[4.25]	[19.23]
H	-868	-1,299	-2,167
	[4.79]	[8.90]	[13.59]
	Rebates		
Feasible	262 -1657	308-1338	570-2995
Observed	1,882	214	2,096
PS	501	-272	229
	[17.39]	[10.75]	[27.53]
CS	261	-110	150
	[11.44]	[6.39]	[16.55]
SS	762	-383	379
	[28.56]	[16.94]	[43.90]

Notes: The top pane reports revenues under three assortments using an e^R e ort policy for each one. The second pane reports changes in variable profit from moving from one assortment to another, as indicated. Rebate ranges in the third pane reflect the IR and IC constraints of the retailer and Mars. Standard errors are computed according to the procedure in Appendix A.6. The reported 'Observed' rebate uses the observed discount in the calculation of the rebate payment. CS assumes a demand elasticity of $\epsilon =$

Table 10: Critical Thresholds and Foreclosure at Observed

$-\frac{MIN}{M}$	$-\frac{MAX}{M}$	Assortment	Effort
0	11,763	(H,M)	$e^R(H;M)$
11,763	11,912	(H,M)	$e^{-M}(H;M)$
11,912	13,101	(M,M)	$e^R(M;M)$
13,101	13,319	(M,M)	$e^{-M}(M;M)$
13,320	1	(H,H)	$e^{NR}(H;H)$

Notes: Calculations report the retailer's optimal assortment and effort policy at the observed for different values of the threshold.

Table 12: Potential Gains from E ort

	Vertically Integrated			Socially Optimal		
	(H,H)	(H,M)	(M,M)	(H,H)	(H,M)	(M,M)
% ($e^{NR}; e^{Opt}$)	9.89	8.61	7.95	13.69	13.11	13.26
% ($e^R; e^{Opt}$)	7.78	6.51	6.18	11.67	11.11	11.58
R	-83	-63	-55	-163	-152	-157
	[2.75]	[2.51]	[2.30]	[4.23]	[3.77]	[3.87]
M	195	152	128	251	211	190
	[5.83]	[5.10]	[4.92]	[6.61]	[5.62]	[5.70]
PS	76	65	63	39	24	17
	[3.09]	[2.65]	[3.04]	[3.32]	[3.30]	[3.64]
$CS(= 2)$	228	210	192	289	290	284
	[5.74]	[5.68]	[5.88]	[5.93]	[5.65]	[6.08]
SS	304	275	255	329	313	301
	[8.51]	[8.02]	[8.63]	[8.45]	[7.78]	[8.54]

Notes: Percentage change in policy is calculated as increase required from baseline policy e^{NR} to vertically integrated or socially optimal policy. Social optimum assumes corresponding to a median own price elasticity of demand of $= 2$. For robustness, see Appendix A.4.

Table 13: Net Effect of Efficiency and Foreclosure

Base: to ($M;M$) with e ort:	$(H;H)$ and e^{NR}			$(H;M)$ and e^{NR}		
	e^R	e^{VI}	e^{SOC}	e^R	e^{VI}	e^{SOC}
R	-575	-626	-728	-312	-364	-466
	[2.39]	[2.75]	[4.21]	[0.93]	[2.31]	[3.79]
M	3,045	3,140	3,201	1,382	1,476	1,538
	[21.59]	[21.74]	[22.07]	[4.88]	[6.05]	[6.72]
PS	267	302	255	-239	-203	-250
	[27.84]	[27.58]	[27.03]	[10.71]	[10.30]	[9.81]
$CS(= 2)$	211	352	444	-49	92	185
	[17.14]	[18.07]	[19.32]	[7.03]	[7.92]	[8.95]
SS	477	654	700	-287	-111	-65
	[44.86]	[45.38]	[46.03]	[17.43]	[17.63]	[18.20]

Notes: Consumer Surplus calibrates to median own price elasticity of $= 2$. Calibration only affects the scale of consumer surplus calculations, not the ranking of various options. For more details see Appendix A.4. Only one of our 1000 bootstrap iterations (SS for the $e^{SOC}(H;M)$ case) yields a different sign than those reported in the table.

Table 14: Potential Upstream Deviations

Base: to $(M;M)$ with e ort:	$(H;H)$ and e^{NR} e^R	$(H;M)$ and e^{NR}
---------------------------------	-------------------------------	----------------------

Table 16: Comparison under Alternate Ownership Structures

	No Merger	M-H Merger	M-N Merger	H-N Merger
AUD Assortment Alternative	$e^{VI}(M;M)$ $e^{NR}(H;H)$	$e^{VI}(H;M)$ $e^{NR}(N;N)$	$e^{VI}(M;M)$ $e^{NR}(H;H)$	$e^{VI}(M;M)$ $e^{NR}(H;H)$
<i>R</i>	-626 [4.11]	-253 [6.15]	-616 [3.80]	-626 [4.11]
<i>M</i>	3,140 [22.10]	2,962 [15.25]	3,091 [20.90]	3,140 [22.10]
<i>Rival</i>	-2,173 [13.60]	-1,458 [1.47]	-2,173 [13.60]	-2,212 [12.75]
<i>M</i>	2,111 [4.71]	2,105 [3.44]	2,309 [4.45]	2,111 [4.71]
<i>PS</i>	302 [27.48]	1251 [12.35]	302 [27.61]	302 [27.48]
<i>CS</i> (= 2)	352 [18.98]	769 [9.47]	337 [18.95]	352 [18.98]
Price to Avoid Foreclosure	13.53 [0.22]	-11.90 [0.15]	9.44 [0.25]	14.04 [0.21]
% Reduction in Rebate ($c = 0.15$)	3.55 [0.51]	42.33 [0.26]	12.24 [0.46]	2.34 [0.49]

Notes: Table compares the welfare impacts of an exclusive Mars stocking policy under alternative ownership structures. This assumes threshold is set at the vertically-integrated effort level.

Appendix

A.1: Proof of Theorems

Proof of Theorem 1:

Note: We can relate our (linear) delta operators to one another via:

$$\Delta = \Delta_M + \Delta_H$$

(A3) provides that $\Delta(M;M) > \Delta(H;H)$. (B3) provides $\Delta(H;M) > \Delta(H;H)$ and (C4) provides that $\Delta(H;M) > \Delta(M;M)$. Thus $\Delta(H;M) > \Delta(M;M) > \Delta(H;H)$.

Absent transfers, if R selects the assortment then $R(H;H) > R(H;M) > R(M;M)$ implies that the equilibrium assortment will be $(H;H)$. If we temporarily ignore $(H;M)$ then (A1)-(A3) say that in a choice between $(M;M)$ and $(H;H)$ it is possible to design a transfer T which leads to assortment $(H;H) \succ (M;M)$ in equilibrium. Likewise, if we temporarily ignore $(M;M)$, then under (B1)-(B3) it is possible to design a transfer that leads to assortment $(H;H) \succ (H;M)$.

A.2: 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 d , can be mapped into \bar{M} (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 $\bar{M} = \max\{0, \frac{M - \bar{M}}{M}\}$. This implies $T = (\bar{M})^2 M$, so that as the threshold increases M is limited in how much surplus he 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, $\bar{e} = e^R$ 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 $d = w_m - c_m$ such that M earns no profit on the marginal unit, and some \bar{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 $(w_m; d; \bar{q}_m)$ as:

$$T = \begin{cases} (p_m - w_m + d) q_m & \text{if } q_m \geq \bar{q}_m \\ (p_m - w_m) q_m & \text{if } q_m < \bar{q}_m \end{cases}$$

One can increase the wholesale price w_m by one unit, and the generosity of the rebate (d) by one unit. Continuing with this procedure, the retailer profits when the threshold is met. For any $q_m \geq \bar{q}_m$ as:

$$T \leq M$$

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

$$\begin{aligned}
 & \begin{cases} < R(a;e) + \alpha M(a;e) - T & \text{if } 2PT \\ > R(a;e) & \text{o.w.} \end{cases}
 \end{aligned}$$

We define

In the case where the rebate is paid, we can express the retailer's problem as:

$$e_1 = \arg \max_e R(e) - c(e) + M(e) \quad \text{s.t.} \quad M(e) = \bar{M}$$

The solution to the constrained problem is given by:

$$e_1 = \max_e R(e) - c(e) \quad \text{where } e \text{ solves } M(e) = \bar{M}$$

If the rebate is not paid then:

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

The retailer's IC constraint:

$$R(e_1) - c(e_1) + M(e_1) \geq R(e_0) - c(e_0) \quad (\text{IC})$$

and the dominant firm's IR constraint:

$$(1 - \alpha) M(e_1) \geq M(e_0) \quad (\text{IRM})$$

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

$$R(e_1) - c(e_1) + M(e_1) \geq R(e_0) - c(e_0) + M(e_0)$$

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

A.4: 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, $Q_t^1 = \sum_{j \in A \cap F} q_{jt}^1$ 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 $q_{jt} = q_{jt}^1 - E[q_{jt}^0]$, the treatment effect of removing product(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 rm during that period.

Table 17: Summary of Sales and Revenues for Four Clusters of Machines

	Group Size	Vends/Visit		Revenue/Visit		Avg Sales/Day	
		Mean	S.D.	Mean	S.D.	Mean	S.D.
A	4	39.0	26.1	28.3	18.7	5.8	1.4
B	7	88.9	39.5	70.6	33.4	24.9	3.0
C	27	56.9	31.5	41.5	23.2	9.2	1.4
D	28	71.6	33.8	54.3	26.8	15.1	2.0

Notes: The 66 machines in our analyses are divided into four groups of machines based on the arrival rate and the amount of revenue collected at a service visit, using a k-means clustering algorithm. Our counterfactual analyses are based on cluster D.

Figure 5: Histogram of Daily Sales for Machines in Group D



Notes: The 28 machines in group D form the basis for our counterfactual exercises. Means and standard deviations for all machine groups are reported in table 17.

5. Fit a Chebyshev Polynomial (order 10) to the average of each computed sequence of profits: $u^R(x; a; j^{\wedge b}); u^M(x; a; j^{\wedge b})$, etc.
6. For every possible value of e use (11) to compute: $(a; e; j^{\wedge b}) = (I - P(e))^{-1} \pi(x; a; j^{\wedge b})$.
7. Use $(a; e; j^{\wedge b})$ to calculate the optimal policies for different groups of agents ($e^{NR}; e^R; e^{VI}; e^{SOC}$) for every a .
8. Compute all of the profit differences $R; M; H$ for Tables 9-15.
9. 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 MLE estimates of the demand parameters (as reported in Table 7). 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.

A.7: 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 :

$$CS(a; x; j) = \frac{1}{NS} \sum_{s=1}^S \log \sum_{j \in a(x^s)} \exp[j + \beta_{ij}(x^s)]^A$$

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

$$CS(a; e) = (I - P(x(e)))^{-1} CS(a; x; j)$$

The remaining challenge is that $CS(a; e)$ relates to arbitrary units of consumer utility, rather than dollars. Recall our utility specification from (4), with $\beta = [\beta; \beta; \beta]$:

$$u_{ijt}(x) = \beta_j + \beta_{jt} + \beta_t + \sum_l \beta_{llt} x_{jl} + \beta_{ijt}$$

Without observable, within-product variation in price, $p_{jt} = p_j$, and ϵ_{jt} is not separately identified from the product fixed effect δ_j . If ϵ_{jt} were identified, then we could simply write $CS(a;e) = \int CS_j(a;e) f_j(p) dp$. Instead, we can calibrate ϵ_{jt} given an own price elasticity:

$$\epsilon_{jt} = \frac{p_{jt}}{s_{jt}} \frac{\partial s_{jt}}{\partial p} = \frac{p_{jt}}{s_{jt}} \int \frac{\partial s_{ij}}{\partial p} f(i|j) d_i = \frac{p_{jt}}{s_{jt}} \int (1 - s_{ij}(i; j)) s_{ij}(i; j) f(i|j) d_i$$

The term ϵ_{jt} does not depend directly on δ_j once we have controlled for the fixed effect δ_j . Thus, we can calibrate own-price elasticities. As is conventional in the literature, we work with the median own-price elasticity, $\bar{\epsilon}(j) = \text{median}_j(\epsilon_{jt}(j))$, and recover δ_j as $\delta_j = \int CS_j(a;e) f_j(p) dp$.

Table 18: Socially Optimal Export Policies (under various elasticities)

	= 1			= 2			= 4		
e^{SOC}	220	224	222	227	232	229	233	238	235
% ($e^{NR}; e^{SOC}$)	16.35	16.10	15.91	13.69	13.11	13.26	11.41	10.86	10.98
% ($e^R; e^{SOC}$)	14.40	14.18	14.29	11.67	11.11	11.58	9.34	8.81	9.27
R	-238	-234	-230	-163	-152	-157	-112	-102	-106
M	285	242	213	251	211	190	219	183	166
PS	-12	-35	-36	39	24	17	66	51	46
CS	645	659	637	289	290	284	128	126	124
SS	633	624	601	329	313	301	193	178	170

Table 19: E ort Decisions of Joint Retailer-Consumer

	= 1			= 2			= 4		
e^{NR}	225	228	226	236	239	237	245	249	247
e^R	224	227	225	234	237	235	242	246	244
e^{VI}	219	223	221	225	230	229	230	236	234
e^{IND}	220	224	222	227	232	229	233	238	235
e^{SOC}	220	224	222	227	232	229	233	238	235

Notes: Reports e ort policies that maximize the combined retailer-consumer surplus, under di erent assumptions for median own-price elasticity when calculating consumer surplus.

Table 20: Socially Optimal E ort Policies (Joint Retailer-Consumer)

	= 1			= 2			= 4		
% ($e^{NR}; e^{OPT}$)	2.22	1.75	1.77	3.81	2.93	3.38	4.90	4.42	4.86
% ($e^R; e^{OPT}$)	1.79	1.32	1.33	2.99	2.11	2.55	3.72	3.25	3.69
R	-10	-6	-7	-19	-13	-16	-29	-23	-26
M	23	14	13	50	33	33	77	60	59
PS	7	5	4	19	13	13	31	26	27
CS	46	37	38	54	44	49	43	41	44
SS	53	42	42	73	57	62	75	67	71

Notes: Reports potential gains realized when e ort is chosen to maximize combined retailer-consumer surplus, under di erent assumptions for median own-price elasticity when calculating consumer surplus.

in our base scenario.

In Table 21, we calculate the optimal assortment decision of a joint Retailer-Consumer pair. We nd that the assortment choice depends on how much weight the retailer places on consumer surplus, or how elastic consumers are. Assuming the retailer places full weight on consumer surplus, at a median own price elasticity of $\epsilon = 2$ the retailer is more or less indi erent between the $(H;M)$ assortment and the $(H;H)$ assortment. As consumers become more elastic, the retailer-consumer pair prefers $(H;H)$, and as they become less elastic the retailer-consumer pair prefers the consumer-optimal assortment $(H;M)$.

We combine foreclosure and e ciency e ects where we treat the retailer-consumer as a jointly maximizing pair in Table 22. When consumers are su ciently inelastic, and the retailer accounts for consumer utility when choosing the assortment, he selects $(H;M)$. In this world, any rebate which induces a switch to $(M;M)$ decreased both producer and consumer surplus. As consumers

Table 22: Joint Retailer-Consumer Net Foreclosure/Efficiency Effect

	$\epsilon = 1$	$\epsilon = 2$	$\epsilon = 2$	$\epsilon = 4$	$\epsilon = 4$
From	$e^{NR}(H;M)$	$e^{NR}(H;M)$	$e^{NR}(H;H)$	$e^{NR}(H;M)$	$e^{NR}(H;H)$
To	$e^{VI}(M;M)$	$e^{VI}(M;M)$	$e^{VI}(M;M)$	$e^{VI}(M;M)$	$e^{VI}(M;M)$
<i>R</i>	-329	-348	-658	-357	-654
<i>M</i>	1326	1345	3019	1368	3064
<i>H</i>	-1280	-1285	-2151	-1290	-2160
<i>PS</i>	-286	-293	177	-287	215
<i>CS</i>	-203	-81	230	-27	137
<i>SS</i>	-490	-374	407	-313	351

Notes: Reports changes under different assumptions for median own-price elasticity when calculating consumer surplus.

Table 23: Profits under Alternate Product Assortments and Stocking Policies

Policy	<i>R</i>	<i>M</i>	<i>M</i>	<i>H</i>	<i>N</i>	<i>R+M</i>	<i>PS</i>	<i>CS</i>
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
e^{NR} (267)	36,399 [24.1]	1,875 [4.4]	11,719 [27.4]	1,302 [8.9]	1,260 [3.7]	48,117 [18.3]	50,679 [18.1]	24,861 [138.9]
e^R (261)	36,394 [24.1]	1,882 [4.3]	11,763 [27.0]	1,299 [8.9]	1,257 [3.7]	48,157 [18.8]	50,713 [18.6]	24,923 [139.9]
e^{VI} (244)	36,335 [22.7]	1,899 [4.1]	11,871 [25.9]	1,290 [8.8]	1,249 [3.7]	48,206 [19.5]	50,744 [19.3]	25,071 [142.6]
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
e^{NR} (263)	36,661 [23.1]	1,609 [2.4]	10,055 [14.8]	2,173 [13.6]	1,285 [3.5]	46,716 [14.3]	50,174 [24.6]	24,601 [131.3]
e^R (257)	36,656 [23.0]	1,617 [2.2]	10,106 [14.0]	2,167 [13.6]	1,282 [3.5]	46,762 [14.6]	50,211 [24.7]	24,662 [132.6]
e^{VI} (237)	36,578 [21.8]	1,640 [1.9]	10,251 [11.7]	2,149 [13.5]	1,272 [3.5]	46,829 [15.3]	50,250 [25.1]	24,830 [135.5]
(M,M) Assortment: Three Musketeers and Milkyway								
e^{NR} (264)	36,090 [24.2]	2,091 [4.9]	13,067 [30.9]	0 [0.0]	1,256 [3.8]	49,156 [21.6]	50,412 [18.1]	24,761 [141.4]
e^R (259)	36,086 [24.1]	2,096 [4.9]	13,101 [30.4]	0 [0.0]	1,254 [3.8]	49,187 [22.1]	50,441 [18.6]	24,812 [142.6]
e^{VI} (243)	36,035 [22.8]	2,111 [4.7]	13,195 [29.3]	0 [0.0]	1,246 [3.8]	49,230 [22.6]	50,476 [19.0]	24,953 [145.1]

Notes: Profit numbers represent the long-run expected profit from a 'representative' machine. Rebate payments are assumed to only be paid under an (M;M) assortment; rebate payments under other assortments are reported in light typeface, but are assumed to not be paid to the retailer. The retailer's optimal assortment under each assortment policy is reported in boldface type. The socially-optimal assortment is (H;M); we denote this with boldface type for the PS and CS columns.

Table 24: Profits after Mars-Hershey Merger

Policy	<i>R</i>	<i>M</i>	<i>M</i> + <i>H</i>	<i>N</i>	<i>M</i> + <i>H</i> + <i>R</i>	<i>PS</i>	<i>CS</i>
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers							
e^{NR} (267)	36,399	2,083	13,021	1,260	49,419	50,679	24,861
e^R (262)	36,395	2,089	13,055	1,257	49,451	50,708	24,913
e^{VI} (245)	36,340	2,105	13,155	1,249	49,496	50,745	25,064
(N,N) Assortment: Butterfinger and Crunch							
e^{NR} (257)	36,594	1,631	10,193	2,707	46,787	49,494	24,295
e^R (251)	36,589	1,639	10,246	2,700	46,835	49,535	24,355
e^{VI} (232)	36,514	1,662	10,386	2,681	46,900	49,581	24,512

Notes: Profit numbers represent the long-run expected profit from a 'representative' machine. Rebate payments are assumed to only be paid under an (*H*/*M*) assortment; rebate payments in light typeface are assumed to not be paid to the retailer.

Table 25: Profits after Mars-Nestle Merger

Policy	<i>R</i>	<i>M</i>	<i>M</i> + <i>N</i>	<i>H</i>	<i>M</i> + <i>N</i> + <i>R</i>	<i>PS</i>	<i>CS</i>
Reeses Peanut Butter Cup (H), Three Musketeers (M)							
e^{NR} (267)	36,399	2,077	12,978	1,302	49,377	50,679	24,861
e^R (262)	36,395	2,082	13,013	1,299	49,409	50,708	24,913
e^{VI} (245)	36,340	2,098	13,114	1,290	49,455	50,745	25,064
Reeses Peanut Butter Cup (H), Payday (H)							
e^{NR} (263)	36,661	1,815	11,341	2,173	48,001	50,174	24,601
e^R (257)	36,656	1,822	11,388	2,167	48,045	50,211	24,662
e^{VI} (239)	36,591	1,842	11,511	2,151	48,102	50,253	24,815
Three Musketeers (M), Milkyway (M)							
e^{NR} (264)	36,090	2,292	14,323	0	50,412	50,412	24,761
e^R (259)	36,086	2,297	14,354	0	50,441	50,441	24,812
e^{VI} (244)	36,040	2,310	14,436	0	50,476	50,476	24,946

Notes: Profit numbers represent the long-run expected profit from a 'representative' machine. Rebate payments are assumed to only be paid under an (*M*/*M*) assortment; rebate payments in light typeface are assumed to not be paid to the retailer.

Table 26: Profits after Hershey-Nestle Merger

Policy	<i>R</i>	<i>M</i>	<i>M</i> + <i>H</i>	<i>N</i>	<i>M</i> + <i>H</i> + <i>R</i>	<i>PS</i>	<i>CS</i>
Reeses Peanut Butter Cup (H), Three Musketeers (M)							
e^{NR} (267)	36,399	1,875	11,719	2,562	48,117	50,679	24,861
e^R (261)	36,394	1,882	11,763	2,556	48,157	50,713	24,923
e^{VI} (244)	36,335	1,899	11,871	2,538	48,206	50,744	25,071
Reeses Peanut Butter Cup (H), Payday (H)							
e^{NR} (263)	36,661	1,609	10,055	3,458	46,716	50,174	24,601