

Efficiency and Foreclosure Effects of All-Units Discounts: Empirical Evidence ^{*}

Christopher T. Conlon[†]
Julie Holland Mortimer[‡]

February 15, 2014

Abstract

We study an All-Units Discount, in which a downstream firm pays a linear wholesale price up to a quantity threshold, beyond which a discount applies to all future and previous units. Such contracts, which are common in many industries, potentially have both efficiency and foreclosure effects. We estimate a structural model of demand and retailer effort to quantify the efficiency gains induced by the contracts, and to identify cases in which the contract results in either efficient or inefficient exclusion of competing products. We show how the contract allocates the cost of a stock-out between upstream and downstream firms, and find evidence that the contract induces inefficient exclusion in the confection industry. Finally, we point out that the impact of upstream mergers in many markets is likely to be felt not through the price in the final-goods market, but rather in the wholesale market. We examine the impact of various upstream mergers on the willingness of the dominant firm to offer rebate contracts, and the impact that the rebate contracts have on social welfare.

^{*}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 Scott Hemphill, Michael Riordan, and seminar participants at Harvard Business School, Rochester University, SUNY Stony Brook, Universitat Autònoma de Barcelona, Universitat Pompeu Fabra, and the International Industrial Organization Conference for helpful comments. Financial support for this research was generously provided through NSF grant SES-0617896. Any remaining errors are our own.

[†]Department of Economics, Columbia University, 420 W. 118th St., New York City, NY 10027. email: cconlon@columbia.edu

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

1 Introduction

Manufacturers use a wide variety of vertical arrangements to align retailers’ incentives with their own. These arrangements may induce retailers to provide efficient levels of effort, mitigating downstream moral hazard. However, they may also result in retailer exclusion of upstream competitors. Vertical rebate contracts, often referred to as All-Unit Discounts (AUD’s), have the potential to induce both of these effects, and have recently attracted the interest of anti-trust authorities as the focus of several important anti-trust cases.¹

Understanding the impact of vertical rebates can be challenging. The tension between the efficiency gains from mitigating downstream moral hazard and the welfare losses from the inefficient exclusion of upstream rivals implies that the contracts must be studied empirically in order to gain insight into the relative importance of the two effects. Unfortunately, most such contracts are considered proprietary information by their participating firms, frustrating most efforts to study them empirically. An additional challenge for empirically analyzing the effect of vertical contracts on downstream moral hazard is the difficulty in measuring downstream effort (both for the upstream firm and the researcher).

We address these challenges by examining an AUD rebate contract used by the dominant chocolate candy manufacturer in the U.S., Mars Inc.. With revenues in excess of \$50 billion, Mars is the third largest privately-held company in the United States (after Cargill and Koch Industries). The AUD rebate contract implemented by Mars consists of three main features: a per-unit wholesale price, a per-unit discount, and a retailer-specific quantity target or threshold. Mars’ AUD contract stipulates that if a retailer’s total purchases exceed his quantity target, Mars pays him a lump-sum amount, which 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 single retail vending operator, MarkVend Inc., for whom we are able to collect extremely detailed information on demand, wholesale costs, and contractual terms. The retailer also agreed to run a large-scale field experiment on our behalf, which provides us with additional insight into how the AUD rebate contract might influence the retailer’s decisions.

¹Intel’s use of an AUD was central to several recent cases. In 2009, *AMD vs. Intel* was settled for \$1.25 billion, and the same year the European Commission levied a record fine of €1.06 billion against the chipmaker. In a 2010 *FTC vs. Intel* settlement, Intel agreed to cease the practice of conditioning rebates on exclusivity or on sales of other manufacturer’s products. Similar issues were raised in the European Commission’s 2001 case against Michelin, and *LePage’s v. 3M*. In another recent case, *Z.F. Meritor v. Eaton* (2012), Eaton allegedly used rebates to obtain exclusivity in the downstream heavy-duty truck transmission market. The 3rd Circuit ruled that the contracts in question were a violation of the Sherman and Clayton Acts, as they were *de facto* (and partial) exclusive dealing contracts.

In order to analyze the effect of the AUD contract, we specify a discrete-choice model of consumer demand and a model of retailer behavior, in which the retailer chooses two actions: a set of products to stock and an effort level. We hold retail prices fixed throughout the analysis, consistent with the data and common practice in this industry.² The number of units the retailer can stock for each product is constrained by the capacity of his vending machines, and we interpret retailer effort as the frequency with which the retailer restocks his machines. In order to calculate a retailer’s optimal effort level, we estimate a dynamic restocking model à la Rust (1987), in which the retailer chooses how long to wait between restocking visits. Due to the capacity constraints of a vending machine, the number of unique products the retailer can stock is relatively small. Thus, we estimate the dynamic restocking model for several discrete sets of products, and we assume that the retailer chooses the set of products to stock that maximizes his profits. These features of the retail vending market (i.e., fixed capacities for a discrete number of unique products) make it well-suited to studying the impacts of the AUD contracts, because the retailer’s decisions are discrete and relatively straightforward.³

Identification of our demand and supply-side models benefits from the presence of exogenous variation in retailer stocking decisions that were implemented for us by the retailer in a field experiment. One approach to measuring the impact of effort on profits might be to persuade the retailer to directly manipulate the restocking frequency, but this has some disadvantages. For example, the effects of effort (through decreased stock-out events) are only observed towards the end of each service period, and measuring these effects might prove difficult. Instead, we focus on manipulating the likely results of reduced restocking frequency – by exogenously removing the best-selling Mars products. We find that in the absence of the rebate contracts, Mars bears almost 90% of the cost of stock-out events, as many consumers substitute to competing brands, which often have higher retail margins. The rebate, which effectively lowers the retailer’s wholesale price for Mars products, reduces Mars’ share of the cost of stock-out events to roughly 50%, and the quantity-target aspect of the rebate provides additional motivation for the retailer to set a high service level.

After estimating the models of demand and retailer behavior, we explore the welfare implications of the retailer’s optimal restocking decisions. Mars’ AUD contract is designed to

²By holding retail prices fixed, we do not require an equilibrium model of downstream pricing responses to the AUD contract. In practice, we see almost no pricing variation over time or across products within a category (i.e., all candy bars are priced the same as each other, and this price holds throughout the period of analysis). Over a short-run horizon of about three to five years, the retailer has exclusive contractual rights to service a location, and these terms may also commit him to a pricing structure during that time.

³These features also characterize other industries, such as brick-and-mortar retail and live entertainment.

induce greater retailer effort through more frequent re-stocking. However, when the retailer increases his re-stocking effort under the contract, he re-stocks all products regardless of manufacturer. Demand externalities across products of different upstream firms imply that the retailer’s optimal stocking decision might lead to either over- or under-supply of retailer effort from a welfare perspective. Over some ranges of the re-stocking policy, more frequent re-stocking *reduces* sales of Hershey and Nestle products, because these products no longer benefit from forced substitution when the dominant Mars products sell out. Downstream effort is substitutable in this range. Over another range of the re-stocking policy, all products stock out, and greater re-stocking effort increases sales of all products (including those of Hershey and Nestle). Downstream effort is complementary in this range. We find evidence that the rebate induces greater retailer effort, and that this effort is substitutable across manufacturers in the confections market we study.⁴

Once we have characterized the retailer’s optimal re-stocking policy, we ask whether or not the downstream firm could increase profits by replacing a Mars product with a competitor’s product. We find evidence that the Mars’ AUD forecloses competition in the market we study. Specifically, the retailer can increase profits by substituting a Hershey product for a Mars product, but the threat of losing the rebate discourages him from doing so.

Finally, we note that the impacts of upstream mergers are often felt not through the price in the final-goods market, but rather in the wholesale market. We simulate the impact of various counterfactual upstream mergers on the willingness of the dominant firm to offer rebate contracts, and the impact that the rebate contracts have on social welfare. Interestingly, we find conditions under which an upstream merger of a dominant firm with a close competitor can lead to socially-efficient downstream effort and product assortment. We also find that an upstream merger of two smaller rivals, while it cannot necessarily prevent exclusion, can bid up the price of a downstream firm’s shelf space.

1.1 Relationship to Literature

There is a long tradition of theoretically analyzing the potential efficiency and foreclosure effects of vertical contracts. The literature that explores the efficiency-enhancing aspects of vertical restraints goes back at least to Telser (1960) and the *Downstream Moral Hazard* problem discussed in Chapter 4 of Tirole (1988). Klein and Murphy (1988) show that without vertical restraints, retailers “will have the incentive to use their promotional ef-

⁴We use a calibrated cost of re-stocking based on average wages of drivers and time spent re-stocking each machine.

forts to switch marginal customers to relatively known brands...which possess higher retail margins.” More directly, Deneckere, Marvel, and Peck (1996), and Deneckere, Marvel, and Peck (1997) examine markets with uncertain demand and stock-out events, and show that vertical restraints can induce higher stocking levels that are good for both consumers and manufacturers.⁵

One of the important developments in the theoretical literature 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 industry profits. Much of the subsequent theoretical literature focuses on demonstrating that the *Chicago Critique’s* predictions are a bit special. For example, Aghion and Bolton (1987) show that long-term contracts that require a liquidated damages payment from the downstream firm to the incumbent can result in inefficient exclusion; while Bernheim and Whinston (1998) show that the *Chicago Critique* ignores externalities across buyers, and that once externalities are accounted for, inefficient exclusion is again possible. Later work by Fumagalli and Motta (2006) links exclusion to the degree of competition in the downstream market. While extremely influential with economists, these arguments have (thus far) been less persuasive with the courts than Bork (1978).

Relatedly, a separate theoretical literature has explored the potential anti-competitive effects of vertical arrangements in the context of upfront payments or slotting fees paid by manufacturers to retailers in exchange for limited shelf space (primarily in supermarkets). This literature includes Shaffer (1991a) and Shaffer (1991b) which analyze slotting allowances, RPM, and aggregate rebates to see whether or not they help to facilitate collusion at the retail level. Sudhir and Rao (2006) analyze anti-competitive and efficiency arguments for slotting fees in the supermarket industry. A broader literature has also examined the conditions under which bilateral contracting might lead to exclusion, such as Rasmusen, Ramseyer, and Wiley (1991), Segal and Whinston (2000), and more recently Asker and Bar-Isaac (2014).

Since the *Intel* anti-trust cases, there has been renewed interest in AUD contracts. Chao and Tan (2013) show that AUD and quantity-forcing contracts can be used to exclude

⁵While the empirical literature on the efficiency effects of vertical restraints is limited, primarily by lack of proprietary data on vertical contracts, Shepard (1993) provides an empirical study of downstream moral hazard in the context of the gasoline retail market and Asker (2005) examines the efficiency and foreclosure effects of exclusive dealing. Other studies, such as Lafontaine (1992) and Brickley and Dark (1987), focus on franchise arrangements rather than the types of rebate contracts we examine here.

a capacity-constrained rival, and O’Brien (2013) shows that an AUD may be efficiency enhancing if both upstream and downstream firms face a moral-hazard problem. Prior to the *Intel* case, Kolay, Shaffer, and Ordovery (2004) showed that a menu of AUD contracts can more effectively price discriminate than a menu of two-part tariffs when the retailer has private information about demand.

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 field experiments, and section 4 provides the details for the empirical implementation of the model. Section 5 provides results, and section 6 concludes.

2 Theoretical Framework

In a conventional nonlinear discount contract, the retailer pays a linear price w for the first $\bar{q}$ units of a good, and then pays $w - \Delta$ (for $\Delta > 0$) thereafter. Under an AUD, the discount applies retroactively to all previous units, as well as all additional units, so that retailer cost is $C(q) = wq - \mathbf{1}[q > \bar{q}] \cdot \Delta \cdot q$. Both contracts are shown in Figure 1. The structure of the AUD implies that for some quantities, the retailer makes a lower total payment but receives more total units of the good. This use of a negative marginal cost has lead some to believe that AUD is de facto evidence of anticompetitive behavior.

A possible defense of the AUD contract (also employed by *Intel*) is that it has the potential to be efficiency enhancing if it encourages the retailer to exert costly effort required to sell the good. This effect enters through both features of the contract: (1) the lower marginal price, $w - \Delta$, and (2) the choice of the threshold $\bar{q}$, which triggers a large transfer payment from the manufacturer to the retailer. Much like a two-part tariff, an appropriate choice of $\bar{q}$ can incentivize an efficient level of downstream effort.⁶

We present a simple framework that provides some intuition for our empirical exercises, though in practice we are able to consider more general settings than the one we describe. We consider a single downstream retailer R , a dominant upstream firm M , and two upstream competitors H, N .⁷ The three upstream firms each sell several competing differentiated products. In the first stage, each of the three upstream firms sets a single linear wholesale price per unit for all of their products, (w_m, w_n, w_h) .⁸ In the second stage, the dominant

⁶Related to the potential quantity-forcing effect of the threshold, it is worthwhile to point out that lower retail prices are a non-contractible form of effort that is costly for the retailer to provide, and demand enhancing for the upstream firms.

⁷We think of M as Mars, and H and N as Hershey’s and Nestle.

⁸Though demand may be different for different products sold by the same manufacturer within a product

firm M proposes a nonlinear rebate contract, which consists of a discount and a threshold $(\Delta, \bar{q})$, where $\bar{q}$ represents the total sales across all of M 's products. After observing the wholesale prices and the contract, the retailer chooses a set of products a , and a level of effort e . We assume that the number of unique products R chooses in a is exogenously determined.⁹ Finally, sales are realized $\mathbf{q}(a, e)$, and depend on both the product assortment and the effort level. We assume that the retailer charges consumers a fixed uniform price across all products (independent of manufacturer). While this assumption is restrictive, it accurately depicts the industry we study (and many others), in which competition is over downstream service quality and product mix, rather than retail prices.

We consider a single scalar version of effort e , rather than product- or manufacturer-specific effort. In our application, effort corresponds to how often a retailer restocks, and all products are restocked simultaneously. The benefit of increased effort is that products are more likely to be available when consumers arrive, thus consumers always benefit from more effort. The cost of increased effort is that restocking is a costly activity for the retailer. Thus, the retailer solves:

$$\max_{a, e} \pi_r(a, e) - c(e). \quad (1)$$

When profits of the dominant upstream firm $\pi_m(a, e)$ are increasing in effort, there is an incentive for M to offer contracts to the retailer that enhance effort. We do not make any restrictions as to whether $\pi_h(a, e)$ or $\pi_n(a, e)$ are increasing or decreasing in effort. The demand externalities that arise from the retailer's effort imply that a vertically-integrated firm consisting of (R, M) might set an effort level that is either too high or too low from a social perspective, depending on whether retailer effort is a substitute or a complement to the profits of the upstream competitors. The upstream competitors might either "free-ride" on the enhanced effort that M induces, or enhanced effort may imply "business stealing" from H, N by M . In our empirical work, we focus on distinguishing between these two possibilities and measuring the degree of substitutability or complementarity of retailer effort.

The second choice the retailer makes is the choice of product assortment a . The rebate contract may induce the retailer to stock more products by M and fewer products by H and N . It may also induce the retailer to select products made by H and N that do not compete closely with products made by M . The retailer can compute the optimal effort

category, uniform pricing is a common feature of many markets. For example, manufacturers do not generally charge different prices for different flavors of yogurt, or for orange juice with or without pulp.

⁹For example, the number of unique products might be determined by shelf-space constraints.

level e for each choice of a for a given set of wholesale prices and rebate contract terms. Given the optimal choice of effort, we assume that the retailer chooses the assortment a , that maximizes his profits (inclusive of potential rebate transfers): $\pi_r(a, e(a)) \geq \pi_r(a', e(a'))$ for all $a' \neq a$.

Once we have characterized the retailer's choice of (a, e) for a given set of wholesale prices, we can determine whether or not a particular rebate contract is individually rational for M to offer, and whether a rebate contract that induces (partial or full) exclusion of H or N is individually rational (IR) and incentive compatible (IC) for R . Conditional on a contract $(\Delta, \bar{q})$, we can also ask whether or not H or N would be willing to set a different wholesale price than the one we observe in order to avoid (full or partial) exclusion. Alternatively, if there is no non-negative price at which H or N could avoid exclusion, we can also consider the amount by which M might be able to reduce the discount Δ and still obtain the same a (i.e.: not violate the IR and IC constraints of the retailer).

The solution concept we employ is subgame perfection. We consider one-shot deviations rather than fully endogenizing the first-stage wholesale prices (w_m, w_n, w_h) , because allowing w_m to freely adjust would result in a continuum of equilibria in our game.¹⁰ Subgame perfection implies that the retailer is unable to pre-commit to a higher level of service (or an enhanced presence in retail product assortment) for a given set of contracts, in order to extract a better deal from the upstream firm.

In line with the theoretical literature, we can examine the effect the rebate has on total industry profits $\pi_{ind} = \pi_r + \pi_h + \pi_n + \pi_m$. The formal prediction of the *Chicago Critique* is that exclusion should only be possible when it maximizes industry profits. The intuition is that the retailer could hold an auction in which firms bid for exclusivity. The game-theoretic literature (e.g., Bernheim and Whinston (1998), and Segal and Whinston (2000)) shows that while an exclusive contract may increase bilateral surplus $\pi_r + \pi_m$, externalities outside the contract imply that it need not maximize π_{ind} .

Our paper departs from the *Chicago Critique* in some key ways. First, we add downstream moral hazard and potential efficiency gains, similar to other literatures on vertical

¹⁰To illustrate, consider increasing the wholesale price to $(w_m + \epsilon)$ and the rebate to $(\Delta - \epsilon)$. This results in the same post-rebate wholesale price $(w_m - \Delta)$, and implies the same cost function for the retailer for any quantity in excess of $\bar{q}$. If we kept increasing both the wholesale price w_m and the rebate Δ , in the limit this approaches a quantity-forcing contract with a linear tariff for quantity in excess of $\bar{q}$. For this reason, we do not consider upward deviations of w_m . In practical terms, this may be justified by the ability of retailers to purchase from other channels (or other retailers). In the case of confections, if wholesale prices increased substantially, the retailer could purchase inventory at warehouse clubs like Costco, supermarkets, or even other retailers.

arrangements. Second, we consider differentiated multi-product upstream firms. Thus, the degree of business stealing and competition may vary across the potential sets of products in *a*. Finally, we restrict the retailer to a specific number of products, rather than the “naked exclusion” of Rasmusen, Ramseyer, and Wiley (1991).

2.1 Endogenous Effort

One of the key contributions of our paper is that it considers both pro- and anti-competitive justifications for rebate contracts, and measures empirically which effect dominates. The crucial issue is whether or not stronger incentives for (efficient) downstream effort counter-balance the potential that AUD contracts have to exclude rival manufacturers. In order to compare the two forces, we need to understand how effort endogenously responds to the different contractual forms and product assortments. In most empirical contexts, the econometrician has very little data on the cost of effort. While the previous section considered a general form of retailer effort e , in this section we consider the specific case in which the retailer chooses the restocking frequency.

We consider a retailer who follows a multi-product $(\mathbf{s}, \mathbf{S})$ policy, in which he pays a fixed cost FC and fully restocks (all products) to target inventory $\mathbf{S}$. The challenge is to characterize the critical re-stocking inventory level, $\mathbf{s}$. For modeling the firm’s decision, it is more convenient to work with the number of potential consumer arrivals, which we denote x , rather than $\mathbf{s}$, because in a multi-product setting, $\mathbf{s}$ is multi-dimensional (and may not define a convex set), while x is a scalar. This implies an informational restriction on the retailer: namely, that he observes the number of potential consumers (for example, the number of consumers who walk through the door) but not necessarily the actual inventory levels of each product when making restocking decisions. This closely parallels the problem of the vending operator that we study in our empirical example. That is, he might have information on whether particular days are likely to be busy or not, but he does not observe the actual inventory levels of individual products until he visits the machine and restocks it.

The retailer solves the following dynamic stocking problem, where $u(x)$ denotes the cumulative retailer profits after x potential consumers have arrived. We assume that profits are not collected by the retailer until he restocks (again this parallels our vending machine setting). The retailer’s value function is:

$$V(x) = \max\{u(x) - FC + \beta V(0), \beta E_{x'}[V(x'|x)]\} \quad (2)$$

The problem posed in (2) is similar to the “Tree Cutting Problem” of Stokey, Lucas, and Prescott (1989), which for concave $u(x)$ and increasing $x' \geq x$, admits a monotone policy such that the firm re-stocks if $x \geq x^*$. Given a guess of the optimal policy, we can compute the post-decision transition-probability-matrix $\tilde{P}$ and the post-decision utility $\tilde{u}$ defined as:

$$\tilde{u}(x) = \begin{cases} 0 & \text{if } x < x^* \\ u(x) - FC & \text{if } x \geq x^* \end{cases}$$

This allows us to solve the value function at all states in a single step:

$$V(x, x^*) = (I - \beta \tilde{P}(x^*))^{-1} \tilde{u}(x, x^*) \quad (3)$$

This also enables us to evaluate profits under alternative stocking policies x' , 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 $u(x) + u^m(x)$, which incorporates the profits of the dominant upstream manufacturer. Likewise, we can consider the industry-optimal policy by replacing $u(x)$ with $u(x) + u^m(x) + u^h(x) + u^n(x)$.

To find the optimal policy we iterate between (3) and the policy improvement step:

$$x^* = \min x : u(x) - FC + \beta V(0, x^*) \geq \beta P(x'|x) V(x', x^*) \quad (4)$$

The fixed point $(x^*, V(x, x^*))$ maximizes the long-run average profit of the agent $\pi_0(x^*) V(x, x^*)$ where $\pi_0 \tilde{P} = \pi_0$ is the ergodic distribution corresponding to the post-decision transition matrix. Once we have obtained the long-run average profits for a given policy, we can compare across different product assortments and contractual forms.

2.2 Numerical Example

In order to illustrate the role that an AUD plays in a retailer’s choice of product assortment, this section provides a numerical example in which an AUD contract can be used to achieve exclusion. In order to keep the example simple, we recast the retailer’s dynamic restocking effort as a capacity constraint (i.e., higher capacity is equivalent to greater retailer effort). This allows us to gain insight into the potential of an AUD contract to induce exclusion through the retailer’s choice of product assortment, using a static model with costless retailer effort. If the efficiency gains from an AUD contract are relatively small, this intuition should

mimic the results of the dynamic model of retailer effort.¹¹

The intuition is captured by considering two of the upstream firms, M and H , selling to retailer R . Firm M sells two products (products 1 and 3) and offers an AUD, and Firm H sells a single product (product 2) under a simple linear price schedule. We assume both upstream firms face production costs of zero, and sell their products to the downstream firm at wholesale prices of (w_M, w_H) (i.e., M sells products 1 and 3 at the same wholesale price). We also assume that the downstream firm chooses two products to stock (i.e., $[1,2]$, $[2,3]$, or $[1,3]$), sets a single price p for all products, and faces a per-product capacity constraint r , which is equal for all products. We examine this simple setting because it captures the important trade-offs of the AUD contract that we observe, and also because it mimics the arrangements within the chocolate confections vending market. However, it is worth noting that many other markets may be similarly characterized.¹²

Consumers choose a single product from the pair of products stocked by the downstream firm, or the outside good (product 0). Each consumer has a preference ranking over two of the four possible products (e.g., a consumer’s preference ranking might be $[1,3]$, or $[2,0]$, etc.), which determine his type. No consumer in the market ranks the outside good first. If a consumer’s first-choice product is not available, he chooses his second-choice product. If neither the first- or second-choice product is available, the consumer exits the market. We choose the distribution of consumer types so that demand is strongest for product 1 and weakest for product 3, and we allow for random consumer arrivals. Consumer arrivals are selected for the whole population with replacement, so demand is stochastic.¹³

The choice variables are: w_M , $\bar{q}$ (the threshold for the AUD), and Δ (for firm M), w_H (for firm H), and a (for the downstream firm). We are interested in the conditions under which the AUD prompts the downstream firm to drop product 2 in favor of product 3 (i.e., to contract exclusively with firm M), and the welfare implications of this decision. We generate random arrivals of 100 consumers with replacement, setting p equal to \$1, w_M equal to \$0.40, w_H equal to \$0.20, Δ equal to \$0.15, and the AUD threshold equal to 65. We assume production costs are zero, and we examine the outcome of the model under two different retail capacity constraints: 45 units for each product, or 65 units for each product. We simulate the model 100,000 times to account for differences in outcomes based on the

¹¹For a brief note on the challenges of explicitly solving for the optimal assortment, see section A.1 of the appendix.

¹²For example, many digital markets such as iTunes display limited price variation, both at the retail and wholesale levels, and do not carry all possible products.

¹³Random selection of consumers without replacement implies deterministic demand, although realized sales may vary based on arrival order and capacity.

random ordering of consumers and the stochasticity of demand.

The results of the model are shown in Table 1. Under either capacity, the rebate is paid if and only if the retailer exclusively sells M 's products [1, 3]. Thus, even in the presence of stochastic demand, the AUD can be used to obtain exclusivity. Under either capacity, the profits of the retail firm and firm M are higher under the exclusive arrangement 74%(99.4%) of the time for the retailer under low(high) capacity, and 99.5%(85.4%) of the time for firm M for the low(high) capacity. The retailer prefers the exclusive arrangement less than 1% of the time in the absence of the AUD. Given the nature of demand and the expected profits of R and M , both firms would have an incentive to agree to an AUD contract that led to de-facto exclusivity.

However, the two capacities yield different predictions for total (industry-level) sales. In the low capacity case (first column), total sales under the non-exclusive stocking arrangement [1, 2] exceed sales under the exclusive [1, 3] about 66% of the time, and are lower about 26% of the time. In the high capacity case, the results are flipped, with higher sales under the exclusive arrangement [1, 3] about 66% of the time. We note that prices and costs are the same across products (and the marginal cost of production is zero), so total sales are equivalent to overall industry profits. Thus, for the low-capacity case, the non-exclusive arrangement maximizes industry profits about two-thirds of the time (i.e., exclusion is efficient), while in the high-capacity case the exclusive arrangement maximizes industry profits about two-thirds of the time (i.e., exclusion is inefficient). Thus, whether the AUD induces efficient or inefficient exclusion depends on retail capacity.

2.3 A Brief Comparison with Other Contracts

An important consideration is how the AUD rebate contract compares to other potential contracts. We consider the four most likely alternatives to the AUD: a purely linear wholesale price, a two-part tariff (2PT), a quantity forcing contract (QF), and a quantity discount (QD). We focus primarily on the efficiency aspect, holding fixed the set of products a . This section is expository, and does not present any original theoretical results.

Throughout our analysis we assume that retail prices are fixed. Following the previous section, we consider the problem of the retailer as trading off variable profit $\pi_r(a, e)$ and cost of effort $c(e)$:

$$\max_{a,e} \pi_r(a, e) - c(e)$$

For the purpose of comparison, we note that the vertically-integrated firm M-R would maximize the joint profits of the retailer and the dominant upstream manufacturer:

$$\max_{a,e} \pi_r(a, e) - c(e) + \pi_m(a, e)$$

With probability $p(a, e)$ (which is increasing in R 's effort and the number of M 's products contained in a), M pays R a transfer $t(a, e)$:

$$\max_{a,e} \pi_r(a, e) - c(e) + p(a, e)t(a, e) \tag{5}$$

In the absence of vertical restraints, and holding the product assortment, a , fixed, the retailer sets the value of e too low: $\pi'_r(e) = c'(e)$. The vertically-integrated firm would set $\pi'_m(e) + \pi'_r(e) = c'(e)$, and it is possible to implement the vertically-integrated effort level through the probabilistic transfer payment from M to R if:

$$p'(e)t(e) + p(e)t'(e) = \pi'_m(e)$$

We can now characterize different contracts. The 2PT achieves the integrated level of e under the familiar sell-out contract, in which M charges a fixed fee and sells at marginal cost: $t'(e) = \pi'_m(e)$ with $p(e) = 1$ and $t < 0$.¹⁴ The QD contract can only achieve the integrated level of effort if $t'(e) = \pi'_m(e)$ (i.e., M sells at marginal cost). To illustrate, note that if $\bar{e}$ denotes the level of effort for which $\bar{q}$ is achieved, $t(\bar{e}) = 0$ by the continuity of the QD contract. Thus, the effect of the QD contract comes completely through marginal cost, because the threat of failing to reach the threshold has no impact on retailer profit. The same is true of the linear wholesale price contract. The AUD has a positive value of $t(\bar{e})$, because it is able to leverage all previous sales (rather than only the marginal unit); thus, the threat of not paying the rebate $p'(e)$ has bite. This means the upstream firm need not give up all of her profit on the margin, so that $\pi'_m(e) - t'(e) > 0$.¹⁵ The QF contract allows M to offer a contract that requires the integrated level of effort, through $\bar{q}$. The only difference between the AUD and the QF contract arises from the fact that the AUD allows a linear schedule both before and after $\bar{q}$, which means the AUD is more flexible when there is

¹⁴The challenge of the 2PT is that the upstream firm M must determine the appropriate fixed fee $t(0)$. Kolay, Shaffer, and Ordovery (2004) shows that a menu of AUD contracts may be a more effective tool in price discriminating across retailers than a menu of 2PTs. Of course, an individually-tailored 2PT enables full extraction by M , but is a likely violation of the Robinson-Patman Act.

¹⁵This leads O'Brien (2013) to show that an AUD contract can enhance efficiency under the double moral-hazard problem (when the upstream firm also needs to provide costly effort such as advertising).

uncertainty about downstream demand. In the absence of this uncertainty, the AUD mimics a QF contract.¹⁶

As a comparison, one can also consider a setting where rebate contracts are not allowed and firms are required to offer linear wholesale prices. Even in this restrictive setting, solving for optimal prices is difficult, because the solution depends both on the effort of the downstream retailer, and the endogenous product assortment. The challenge is that small changes in wholesale prices may induce large discrete changes in the assortment. This is the reason we focus on deviations from equilibrium behavior in the confections industry.¹⁷

3 The Vending Industry and Experimental Data

3.1 Vertical Arrangements in the Vending Industry

AUD rebate programs are the most commonly-used vertical arrangement in the vending industry. Under a rebate program, a manufacturer 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, typically expressed as a percentage of year-over-year sales. The sales goal for an operator is typically set for the combined sales of a manufacturer’s products, rather than for individual products. Some manufacturers also require a minimum number of product “facings” in an operator’s machines. The amount of the rebate and the precise threshold of the sales goal or facing requirement is specific to an individual vending operator, and these terms are closely guarded by participants in the industry.

We are fortunate in that we observe the specific terms of the Mars Gold Rebate program; we include some promotional materials in Figure 1. The program employs the slogan *The Only Candy You Need to Stock in Your Machine!*, and provides a list of ‘must-stock’ items (Snickers, M&M Peanut, M&M Plain, Twix, a choice of 3 Musketeers or Milkyway, and a choice of Skittles or Starburst), as well as a sales target (90% of quarterly year-over-year sales) that applies to the total cases of Mars products sold. We also observe, but are not allowed to directly report, the amount of the rebate. Unlike the Intel rebate program, these rebates do not explicitly condition on marketshare or the sales of competitors. However, they do mandate 6 ‘must-stock’ items, and most vending machines typically carry only 6 or 7 candy bars. While there is some ability for the vending operator to adjust the overall number of candy bars in a vending machine, it is often technologically difficult to do without

¹⁶Chao and Tan (2013) explore connections between QF, AUD, and 3PT when a dominant manufacturer faces a capacity-constrained rival.

¹⁷For a brief note on the challenges of solving for endogenous linear prices, see section A.2 of the appendix.

upgrading capital equipment because candy bars and potato chips do not use the same size ‘slots.’

By structuring payments as a rebate, rather than directly reducing wholesale price, manufacturers are able to tailor the amount of the cost reduction to each individual operator, and to match it to targets that are retailer specific (e.g., 90 percent of his previous year’s sales).¹⁸ Kolay, Shaffer, and Ordovery (2004) demonstrate that AUD contracts may be more effective for the upstream firm when discriminating across retailers than a menu of two-part tariffs. Rebates may also have an anti-competitive effect, because one way of increasing sales of a manufacturer’s product is to stock fewer competing products.

3.2 Data Description

All of our price and quantity data are provided by Mark Vend Company. Data on the quantity and price of all products vended are recorded internally at each vending machine used in our experiment. The data track vends and revenues since the last service visit (but do not include time-stamps for each sale). Any given machine can carry roughly 35 products at one time, depending on configuration. We observe prices and variable costs for each product at each service visit during our 38-month panel. There is relatively little price variation within a site, and almost no price variation within a category (e.g., chocolate candy) at a site. Very few “natural” stock-outs occur at our set of machines.¹⁹ Over all sites and months, we observe 185 unique products. We consolidate some products with very low levels of sales using similar products within a category produced by the same manufacturer, until we are left with the 73 ‘products’ that form the basis of the rest of our exercise.²⁰

All of these data are recorded at the level of a service visit to a vending machine. Because machines are serviced on different schedules it is sometimes more convenient to organize observations by machine-week, rather than by visit. When we do this, we assume that sales are distributed uniformly among the business days in a service interval, and assign those to weeks. Because different experimental treatments start on different days of the week, we allow our definition of when weeks start and end to depend on the client site and experiment.²¹

¹⁸Robinson-Patman prevents manufacturers from directly price discriminating across competing downstream firms when selling ‘inputs.’

¹⁹Mark Vend commits to a low level of stock-out events in its service contracts.

²⁰For example, we combine Milky Way Midnight with Milky Way, and Ruffles Original with Ruffles Sour Cream and Cheddar.

²¹At some site-experiment pairs, weeks run Tuesday to Monday, while others run Thursday to Wednesday.

In addition to the data from Mark Vend, we also collect data on the characteristics of each product online and through industry trade sources.²² For each product, we note its manufacturer, as well as the following set of product characteristics: package size, number of servings, and nutritional information.²³

In Table 2 we report the national sales ranks as reported from the industry association, as well as the aggregate shares within the candy category from the Mark Vend data. We also report the percentage of machine-weeks in which the product is stocked by Mark Vend. There are some patterns that emerge. The first is that Mark Vend stocks some of the most popular products sold by Mars Inc. (Snickers, Peanut M&Ms, Twix, and Skittles) in most of the machines in our sample. However, Mark Vend only stocks Hershey’s best-selling product (Reese’s Peanut Butter Cups) in 29% of machine-weeks, and it constitutes less than 4% of candy sales, even though nationally it is the fourth most popular product. Likewise, Nestle’s best-selling product, Butterfinger, represents 2.7% of Mark Vend’s sales and is only stocked about one-third of the time. Hershey’s with Almonds is the tenth most popular candy product nationally, and isn’t stocked at all by our retailer. Milky Way, another Mars product, is somewhat overrepresented in our sample, as are Rasinets, a Nestle product, which is stocked in 78% of machine weeks and constitutes almost 9% of overall sales, despite being ranked 17th nationally.

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 that the product mix is geared towards those tastes. These are mostly high-income, professional office workers in Chicago, and they may have very different tastes than consumers from other demographic groups.²⁴ The alternative is that the rebate contracts may induce the retailer to substitute from Nestle and Hershey brands to Mars Inc. brands when making stocking decisions. Similarly, it might be the case that when the retailer does stock brands from competing manufacturers (e.g., Raisinets), they choose brands that do not steal business from key Mars Inc. brands.

²²For consolidated products, we collect data on product characteristics at the disaggregated level. The characteristics of the consolidated product are computed as the weighted average of the characteristics of the component products, using vends to weight. In many cases, the observable characteristics are identical.

²³Nutritional information includes weight, calories, fat calories, sodium, fiber, sugars, protein, carbohydrates, and cholesterol.

²⁴For example, Skittles, a fruit flavored candy sold by Mars is primarily marketed to younger consumers.

3.3 Experimental Design

In addition to observing Mark Vend’s rebate contracts, we were able to exogenously remove one or two top-selling Mars confection products from a set of 66 vending machines located in office buildings, for which demand was historically quite stable.²⁵

Implementation of each product removal was fairly straightforward; we removed either one or both of the two top-selling Mars, Inc. products from all machines for a period of roughly 2.5 to 3 weeks. The focal products were Snickers and Peanut M&Ms.²⁶ The dates of the interventions range from June 2007 to September 2008, with all removals run during the months of May - October. We collected data for all machines for just over three years, from January of 2006 until February of 2009. During each 2-3 week experimental period, most machines receive service visits about three times. However, the length of service visits varies across machines, with some machines visited more frequently than others.

4 Empirical Analyses

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 , and use a superscript 1 to denote sales when a focal product(s) is removed, and a superscript 0 to denote sales when a focal product(s) is available. We denote the set of available products as A , and F as the set of products we remove. Then $Q_t^1 = \sum_{j \in A \setminus F} q_{jt}^1$ and $Q_s^0 = \sum_{j \in A} q_{js}^0$ are the overall sales during a treatment week, and control week respectively. It is also convenient to write the sales of the removed products $q_{fs}^0 = \sum_{j \in F} q_{js}^0$. 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 .

In principle, this calculation is straightforward. In practice, however, there are two challenges in implementing the removals and interpreting the data generated by them. First, there is considerable 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

²⁵In addition to the three treatments described here, we also ran five 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.

²⁶Whenever 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, the firm wanted to minimize the number of phone calls received in response to the stock-out events.

over long periods of time, because 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 section A.3 of the Appendix.

4.1 Parametric Specifications of Demand

In addition to computing treatment effects, we also specify two parametric models of demand: nested logit and random-coefficients logit, which are estimated from the full dataset (including weeks of observational data that do not meet any of our control criteria).

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

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt}. \quad (6)$$

The parameter δ_{jt} is a product-specific intercept that captures the mean utility of product j in market t , and μ_{ijt} captures individual-specific correlation in tastes for products.

In the case where $(\mu_{ijt} + \varepsilon_{ijt})$ is distributed generalized extreme value, the error terms allow for correlation among products within a pre-specified group, but otherwise assume no correlation. This produces the well-known nested-logit model of McFadden (1978) and Train (2003). In this model, consumers first choose a product category l composed of products g_l , and then choose a specific product j within that group. The resulting choice probability for product j in market t is given by the closed-form expression:

$$p_{jt}(\delta, \lambda, a_t) = \frac{e^{\delta_{jt}/\lambda_l} (\sum_{k \in g_l \cap a_t} e^{\delta_{kt}/\lambda_l})^{\lambda_l - 1}}{\sum_{\forall l} (\sum_{k \in g_l \cap a_t} e^{\delta_{kt}/\lambda_l})^{\lambda_l}} \quad (7)$$

where the parameter λ_l governs within-group correlation, and a_t is the set of available products in market t .²⁷ The random-coefficients logit allows for correlation in tastes across observed product characteristics. This correlation in tastes is captured by allowing the term μ_{ijt} to be distributed according to $f(\mu_{ijt}|\theta)$. A common specification is to allow con-

²⁷Note that this is not the IV regression/‘within-group share’ presentation of the nested-logit model in Berry (1994), in which σ provides a measure of the correlation of choices within a nest. Roughly speaking, in the notation used here, $\lambda = 1$ corresponds to the plain logit, and $(1 - \lambda)$ provides a measure of the ‘correlation’ of choices within a nest (as in McFadden (1978)). The parameter λ is sometimes referred to as the ‘dissimilarity parameter.’

sumers to have independent normally distributed tastes for product characteristics, so that $\mu_{ijt} = \sum_l \sigma_l \nu_{ilt} x_{jl}$ where $\nu_{ilt} \sim N(0, 1)$ and σ_l represents the standard deviation of the heterogeneous taste for product characteristic x_{jl} . The resulting choice probabilities are a mixture over the logit choice probabilities for many different values of μ_{ijt} , shown here:

$$p_{jt}(\delta, \theta, a_t) = \int \frac{e^{\delta_{jt} + \sum_l \sigma_l \nu_{ilt} x_{jl}}}{1 + \sum_{k \in a_t} e^{\delta_{kt} + \sum_l \sigma_l \nu_{ilt} x_{kl}}} f(v_{ilt} | \theta) \quad (8)$$

In both the nested-logit and random-coefficient models, we let $\delta_{jt} = d_j + \xi_t$; that is, for each product-market pair, we specify a product intercept and a market-specific demand shifter. We include an additional ξ_t for each of our 15,256 machine-visits. For the nested-logit model, we allow for heterogeneous tastes across five major product categories or nests: chocolate candy, non-chocolate candy, cookie, salty snack, and other.²⁸ For the random-coefficients specification, we allow for three random coefficients, corresponding to consumer tastes for salt, sugar, and nut content.²⁹

4.2 Dynamic Model of Re-stocking

In order to estimate the dynamic restocking model specified in Section 2.1, we use the following procedure. To obtain $\pi(x)$, we use the demand system generated by the random-coefficients model to simulate consumer arrivals and update inventories accordingly. We use actual data on full machine capacities (which varies very little across machines). We simulate 100,000 chains and construct the expected profit after M consumers arrive. We define our state variable to be the number of consumers expected to make a purchase from a “full machine” containing a superset of products. Our “full machine” is defined to contain all of the products in Table 3 plus all of the products in the confections category.

We recover the transition matrix $P(x'|x)$ to match the observed distribution of incremental daily sales. This is similar to Rust (1987), which uses the observed distribution of

²⁸The vending operator defines categories in the same way. “Other” includes products such as peanuts, fruit snacks, crackers, and granola bars.

²⁹We do not allow for a random coefficient on price because of the relative lack of price variation in the vending machines. We also do not include random coefficients on any discrete variables (such as whether or not a product contains chocolate). As we discuss in Conlon and Mortimer (2013a), the lack of variation in a continuous variable (e.g., price) implies that random coefficients on categorical variables may not be identified when product dummies are included in estimation. We did estimate a number of alternative specifications in which we include random coefficients on other continuous variables, such as carbohydrates, fat, or calories. In general, the additional parameters were not significantly different from zero, and they had no appreciable effect on the results of any prediction exercises.

incremental mileage. After converting the expected profits from a function of the number of consumers, to a function of the number of consumers who would have made a purchase at a hypothetical “full machine” (which contains the set of products described above), we then fit a smooth Chebyshev polynomial, and use this as our approximation of profits, $\hat{\pi}(x)$.³⁰

There is some heterogeneity in both the arrival rate of consumers to machines, as well as the service level of different machines, so we divide our sample into four groups based on the arrival rate, and the amount of revenue collected at a service visit. We use a k-means clustering algorithm and report those results in Table 4. Our counterfactual analyses are based on cluster D, which is the largest cluster, and corresponds to 28/66 machines in our sample. Machines in clusters A and C are smaller in size, while those in cluster B represent the very highest volume machines in the sample. We focus on cluster D because it is a large cluster of ‘higher than average volume’ machines, which we think is the most important determinant of the re-stocking decision of the firm. Figure 4 plots a histogram of daily sales for the machines in cluster D, which determines the transition rule for our re-stocking model.

We choose a daily discount factor $\beta = 0.99981$, which corresponds to a 7% annual interest rate. We calibrate $FC = \$10$. This 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 $FC = \{5, 15\}$, which generate qualitatively similar predictions. In theory, we should be able to estimate FC directly off 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 refilling well before any products have stocked-out. This is helpful as an experimental control, but makes identifying FC from data impossible.³¹

5 Results

We begin by discussing the results of our three exogenous product removals. In the first case we remove Snickers, in the second we remove Peanut M&Ms, and in the third we remove

³⁰We designate our state space in terms of expected sales under a “full machine” rather than the market size, because the share of the outside good is often large in discrete choice demand settings. This needlessly increases the dimension of the state space without any additional information. Also, under the hypothetical “full machine” with outside good share s_0 , the relationship between the number of consumers in the demand system M and the state space x is well defined, because $x \sim \text{Bin}(M, 1 - s_0)$ by construction. In practice this merely requires inflating all of the “inside good” probabilities by $\frac{1}{1-s_0}$ when simulating consumer arrivals to compute $\pi(x)$. The fit of the 10th order Chebyshev polynomial is in excess of $R^2 \geq 0.99$

³¹We do not consider possible dynamic considerations, where a lower service level leads to a lower arrival rate of consumers (i.e., as consumers facing stock-outs grow discouraged and stop visiting the machine, or the client location terminates Mark Vend’s service contract).

both products. These products correspond to the top two sellers in the chocolate candy category, both at Mark Vend and nationwide. They are also the two best-selling brands for Mars Inc. as a whole. We can think of these as the *dominant* brands within the category.

We report the results from the joint removal in table 5, and summarize substitution from the focal (stocked-out) product to the top five substitutes in table 6.³² In the joint removal (shown in table 5), 93 consumers substitute to Reese’s Peanut Butter Cups, which represents an 85.6% increase in sales for the Hershey product. Note that Reese’s Peanut Butter Cups were not available in either of the single-product removals. In that same experiment, nearly 123 consumers substitute to other Assorted Chocolate products within the same product category, representing an increase of 117%. This includes several products from Mars such as Milky Way and Three Musketeers, but also some products from other manufacturers, such as Nestle’s Butterfinger. Meanwhile, Raisinets (Nestle), a product that Mark Vend stocks very frequently compared to national averages, sees an increase in sales of only 17% when both products are removed.³³ This gives some indication that Raisinets is not a close competitor to Snickers, and competes less closely with Mars Inc. products than other products within the chocolate candy category.

Table 6 shows that in general, the substitution patterns we recover are reasonable; the top substitutes generally include Snickers or Peanut M&Ms if only one of the products is available. Twix, the third-best selling Mars Inc. brand both nationally and in our sample, is also a top substitute. Consumers also substitute to products outside the chocolate candy category, such as Planters Peanuts or Rold Gold Pretzels.

5.1 Substitution Patterns and Profits

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 Inc., several of them are not. If those other brands have similar (or higher) margins for Mark Vend, substitution may cause the costs of the product removal to be distributed unevenly across the supply chain. Table 7 summarizes the impact of the experiments on Mark Vend, our retailer. In the absence of any rebate agreements, we see the following results. Total vends go down 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

³²Detailed results from the two single removals are described in Conlon and Mortimer (2010), and are available from the authors upon request.

³³Substitution to Raisinets is only 3.3% when Snickers is removed by itself.

only by \$10.74 (a statistically insignificant decline). Similarly, in the joint product removal, overall vends decline by 282.66 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).³⁴

Table 8 examines the impact of the product removals on the upstream firms. Removing Peanut M&Ms costs Mars 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. 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 M&M Peanut). We also see the impact of each product removal on other manufacturers. Hershey (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. Likewise, we see slightly more substitution to the two Nestle products in the Snickers removal, so that Nestle gains \$19.32 (as consumers substitute to Butterfinger and Raisinets); however, Nestle’s gains are a smaller percentage of Mars’ losses in the other two removals.

Finally, we examine the potential efficiency impact of the rebate. There are two ways to understand this. The first is that, similar to a two-part tariff, the rebate lowers the marginal cost to the retailer (and reduces the margin of the manufacturer). This helps to better incentivize the retailer to maintain a high level of availability. 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. Under the rebate contract, the retailer now bears about 50% of the cost of the Peanut M&Ms removal, 40.5% of the cost of the joint removal, and the majority of the cost of the Snickers removal. 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 the overall service level.

5.2 Full Model Results

We now consider the results from the full model, in which we allow the stocking decisions to endogenously vary with the rebate contract and the product mix. For this scenario,

³⁴One reason that total losses appear smaller in the double-product removal is the smaller sample size of viable treatment weeks (89) as compared to 115 for the Peanut M&Ms removal.

we estimate random-coefficients and nested-logit models of demand, and simulate consumer arrivals for the representative product mix in table 3. The parameter estimates of the random-coefficient and nested-logit models are reported in table 9. For our representative machine, we fix the five most commonly-stocked confections products: four Mars products (Snickers, Peanut M&Ms, Twix, and M&M Plain), and Nestle’s Raisinets. We then allow the retailer to consider different choices for the final two slots in the confections category. The potential additional products include three Mars products (Skittles, Milkyway, and 3 Musketeers), two Hershey products (Reese’s Peanut Butter Cup, and PayDay), and two Nestle products (Butterfinger and Crunch).³⁵ Throughout this exercise, we assume that all candy manufacturers have the same manufacturing costs of $c = 0$.

In table 10, we report the optimal restocking policies from our model of endogenous stocking decisions. The first column reports the optimal restocking policy that would be set by the retailer in the absence of any vertical contracts, the second column reports the policy the retailer would set at post-rebate prices, and the third corresponds to the policy the retailer would set if he fully internalized the impact of his decision on Mars (but not other upstream firms). The policies are reported in terms of the state variable x , which denotes expected sales at a full machine (which has capacity of around 380-400). Thus, a larger number corresponds to waiting longer between service visits. In the first row (in which the retailer stocks Reeses and 3 Musketeers), the vertically-integrated firm would choose to service after 246 consumers, while the retailer would service after 267 consumers in the absence of any rebates or incentives. The difference is due to the fact that the retailer pays the full cost of restocking, but because of substitution, the upstream firm bears the majority of the cost of the stockout. In general, the retailer chooses a service level (reported in the fourth column) that is between 8-10% too low, compared to the vertically-integrated firm. The rebate induces faster restocking: in the first row, the retailer now services after 262 sales. However, in general, the price impact of the rebate ($\Delta\pi'(s)$ in the restocking model) only increases the service level by 1.6-2.5%, despite the fact that it more evenly allocates the cost of the stockout across the upstream and downstream firms. The preferred policy of the integrated firm depends on the product mix. When the two additional products are Nestle products (Butterfinger and Crunch) the vertically-integrated firm would set the highest stocking level (replacing after 233 consumers), and the gap between the retailer and

³⁵We do not have sufficient information on other products to consider them in our counterfactual analysis. For example, Hershey with Almonds is popular nationally, but is rarely stocked in our data. As a robustness test, we also consider substituting the five base confection products, but the retailer is always worse off in this case, and for space concerns we do not report those results.

the vertically-integrated firm's incentives are largest, at about 10%. When both additional products are owned by Mars (3 Musketeers and Milkyway) the difference in incentives is smallest, at 8.2%.

An important implication of the stocking policies in table 10 relates to the demand externalities of the retailer's effort across the different upstream firms. Figure 5 reports the expected variable profits per consumer as a function of the restocking policy, using the fitted (Chebyshev polynomial-based) profit function. The figure reports results for the retailer, Mars, Hershey, and Nestle; the cost of restocking is not shown. In the first subfigure, for the retailer, we see that substitution between products of all the manufacturers results in continuously-declining variable profit as restocking is delayed. For Mars, which owns the dominant brands and several other products in the machine, the same effect holds: the longer the retailer delays restocking, the lower is Mars' expected profit per consumer. For Hershey and Nestle, however, we see a different profile. Over an initial range of restocking policies (below 200), per-consumer profits are roughly steady, as loyal consumers purchase Hershey and Nestle products from well-stocked machines. For retail stocking policies that range from about 200 to 450, however, Hershey and Nestle's per-consumer profits increase. This results from forced substitution: as Mars' dominant brands stock-out, Hershey and Nestle sell more units because consumers substitute to their products. Over this range, downstream effort is substitutable across upstream firms: Hershey and Nestle benefit when retailer effort is low because their products substitute for retailer restocking effort. For stocking policies that wait for more than 450 sales, Hershey and Nestle products also stock out, and their profits decline. Over this range, downstream effort is complementary across upstream firms: Hershey and Nestle can free-ride off of Mars' payments to induce greater retailer effort. The range of optimal restocking policies identified in table 10 is 233 to 267, identifying this market as one in which downstream effort is substitutable across upstream firms. This leads to the possibility that a vertically-integrated (Mars-retailer) firm will set a higher effort level than the social (or industry-wide) optimum, in order to avoid substitution to Hershey or Nestle products. Indeed, table 10 confirms that the industry-wide optimal restocking delays longer than the vertically-integrated restocking.³⁶

Table 11 reports profit impacts for alternative product assortments and stocking policies. Each panel denotes a different product mix, and each row corresponds to a policy (retailer, rebate, or integrated) from table 10.³⁷ We implicitly assume that Mars chooses the rebate

³⁶For example, in the first row, the industry would wait for 249 sales before restocking, but the vertically-integrated Mars-retailer firm would wait for 246 sales.

³⁷We display the three product mixes that result in optimal policies under at least one scenario (i.e.,

threshold so that it only pays the rebate if total sales exceed a level that requires both Mars products to occupy the last two slots. Under the vertically-integrated optimum policy, we assume that Mars is able to set the rebate threshold exactly so that the retailer optimizes joint profits.

We find that in the absence of the rebate contract, the retailer would choose the retailer optimal stocking policy, restocking after 263 sales, and carrying Reeses and Payday in the last two slots, which we denote (H, H) , for a profit of \$36,668. However, under the observed rebate contracts, the retailer would prefer instead to stock the two Mars products (3 Musketeers and Milkyway) which we denote (M, M) , and receive a rebate payment. Under the rebate contract, at post-rebate prices, the retailer re-stocks after 261 sales and earns \$36,099 in profits, and another \$2,100 in rebate payments. Mars earns \$13,123-\$2,100= \$11,023 after paying the rebate. Thus, the rebate contract increases the bilateral profits of both Mars and the retailer. We can verify this by examining the difference in the profits of the vertically integrated (Mars + Retailer) firm, which is \$49,222 with the rebate, and \$46,759 without the rebate. These bilateral gains come primarily at the expense of Hershey, which loses \$2,173. The *Chicago Critique* gives the prediction that we should only see exclusion if it increases the overall industry profits (Mars + Retailer + Hershey + Nestle), otherwise Hershey and/or Nestle would be able to prevent exclusion. When we compute that profit difference we find that the rebate increases industry profits from \$50,216 to \$50,509 or by \$263, thus the exclusion appears to be efficient. We also find that the vertically-integrated firm would choose (M, M) , and would set an even higher service level than under the rebate (servicing after 245 sales), with a profit of \$49,260 or \$50,509 at the industry level.³⁸

While the rebate contract increases both bilateral surplus and industry profits as the product mix moves from (H, H) to (M, M) , this assortment does not maximize overall industry profits (assuming the retailer selects the restocking policy in this case). The overall industry profits are maximized under (3 Musketeers and Reeses' Peanut Butter Cup) or (M, H) . Under the retailer optimal stocking policy, this implies an industry profit of \$50,726, but the retailer only receives \$36,398: lower than either (H, H) or (M, M) .³⁹

Retailer Optimal, Rebate Optimal, Vertically-Integrated Optimal, or Industry Optimal). The remaining four product mixes from table 10 are suppressed to make the table easier to read.

³⁸The next-best set of products to stock would be Reeses/Butterfinger, with similar results, but with the impact of the rebate spread more evenly across Hershey and Nestle.

³⁹We do not bold the policies for the industry-optimal results, because although the industry-optimal is an important benchmark for understanding the rebate contracts, unlike the other three policies (Retailer, Rebate, and Integrated), there is not a coordinating mechanism by which this outcome could actually be implemented.

Table 12 considers the impact of a Mars-Hershey merger. We assume that the Hershey product (Reeses) is priced at the Mars wholesale price and included in Mars' rebate contract. Interestingly, the merger alleviates some of the tension between the substitutability of downstream effort across upstream firms, because now the merged Mars-Hershey firm is happy for consumers to substitute to Reeses. This results in an alignment between the industry-optimal set of products and both the rebate and integrated outcomes: the Mars-Hershey merger results in the industry-optimal outcome, either through the use of the AUD rebate, or by integrating with the retailer. In the absence of the rebates, however, the retailer will choose to stock two Nestle products (Butterfinger and Crunch), with lower industry profit and consumer surplus.

Table 13 considers the impact of a Mars-Nestle merger. Again, we set the wholesale price of the Nestle products to Mars' wholesale price after the merger. Substitutability between Mars' products and Nestle's products (namely Butterfinger and Crunch) is much lower than between Reeses and Mars' products, and unlike the Mars-Hershey merger, we no longer achieve the industry-optimal choice of products in the retail market. We essentially replicate the patterns from table 11.⁴⁰ Table 14 considers the impact of a Hershey-Nestle merger, with similar results.

In table 6, we compare the exclusive (M, M) assortment to several alternatives. In the first two columns, we compare (M, M) to two alternatives: the (H, H) assortment from table 11, and an alternative assortment of Reeses PB Cup and Butterfinger from table 10, which we denote (H, N) . In the third column, we compare the optimal assortment from the Mars-Hershey merger in table 12 to the assortment preferred by the retailer in the absence of a rebate, which we denote (N, N) . In the last three columns, we provide similar comparisons under the Mars-Nestle and Hershey-Nestle mergers.

In the first column, we report the retail profit (with and without rebate payments) under (H, H) . The retailer earns \$1,497 from the rebate. Hershey's loses \$2,173, for a combined effect of -\$676. Is this inefficient exclusion? If Hershey lowers its wholesale price to 13.3 cents, the retailer will choose to stock (H, H) , and an outcome in which Mars is exclusive will not be viable. If Hershey's production costs are above 13.3 cents, however, they will be unable to avoid being excluded from the market. By implementing the rebate, Mars increases its profits by \$1,004, and industry profits increase by \$293. Consumer surplus increases both because of the preferred (M, M) product selection, and because the rebate induces faster restocking

⁴⁰Note that unlike table 12, the (N, N) outcome is not attractive to the retailer in this case, because the wholesale price of Nestle products has been set to the Mars' wholesale price in this exercise.

(which we denote as Δ CS (Efficiency)). The rebate, however, fails to induce the socially (or industry) optimal choice of (H, M) , leading to lower industry profits and consumer surplus of \$260 and \$340 respectively. In the second column, we replicate the same exercise, but we compare a Mars exclusive stocking policy to a (H, N) policy in which Reeses and Butterfinger are stocked. In this case, the retailer gains \$1,516 from the rebate, but the loss to the largest single competitor (Hershey) is only \$1,316. The combined change (\$200) implies that there is no wholesale price at which Hershey can prevent exclusion. In fact, Mars can now induce exclusion with a lower rebate payment. Even after lowering the rebate payment by \$200, or 9.5% of the \$2,100 rebate payment, Mars maintains exclusivity. Again, industry profits, Mars' profits, and consumer surplus are higher, and we realize efficiency gains from the rebate due to the greater downstream stocking effort. However, the exclusive arrangement with Mars still falls short of the socially-optimal outcome.

In the third column, we investigate the impact of the rebate contract after a Mars-Hershey merger. The rebate now induces the retailer to take the (H, M) product assortment of Reeses and Three Musketeers instead of his optimal choice of (N, N) . This results in retail profits that are \$1,579 higher under the rebate contract. Nestle loses \$1,455, so there is no positive wholesale price under which Nestle can prevent exclusion. Indeed, the combined Mars-Hershey firm could reduce rebate payment to the retailer by 4.6% and still maintain exclusivity. Industry profits and consumer surplus all increase, and we obtain the industry-optimal product assortment.

In the fourth, fifth, and sixth columns, we repeat the exercise in a counterfactual situation in which Mars merges with Nestle, or Nestle merges with Hershey. Under a Mars-Nestle merger, Hershey could avoid exclusion by lower its wholesale price to 8.3 cents, but this is quite low relative to production costs (which are probably in the range of 15-22 cents). In contrast, the last two columns show that a combined Hershey-Nestle firm could avoid exclusion at a wholesale price of roughly 24 cents. However, avoiding exclusion would be bad for the industry, and would likely result in Mars increasing their rebate. In this sense, most of the gains of a Hershey-Nestle merger go to the retailer, by essentially bidding up the price that Mars pays for the retailer's valuable shelf space.

6 Conclusion

We examine confection sales in the vending machine industry and conduct a field experiment that simulates out-of-stock events. We find that in the absence of rebate contracts, the upstream firm bears approximately 80-90% of the costs associated with stock-out events, and

that rebates more evenly allocate the cost of these events across upstream and downstream firms. We also find that the rebate contracts can be used by Mars, the dominant manufacturer, to prevent the retailer from stocking competing brands owned by Hershey and Nestle. We find that the vertical contracts are rational for both Mars and for the retailer, but do not maximize overall industry profits, in contrast to the *Chicago Critique*.

This is potentially important for several reasons, one of which is that product assortment has been the subject of many recent debates about obesity and the appropriate public policy response to the mix of products offered in vending machines, particularly in school settings.⁴¹ The potential exists for rebate contracts to serve not just as barriers for existing competitors, but also for potential entrants producing healthier products.

⁴¹Forty states now tax junk food or soda products, and cities, school districts, and other local jurisdictions have proposed or implemented restrictions on the set of products that may be offered in vending machines. See Engber (2009) for a recent press article summarizing many policy responses in this area. More recent examples include rules requiring that the mix of beverages in city vending machines favor water in New York City, a ban on sales of sugary drinks in city buildings in San Francisco, and a similar proposed ban in Boston (Smith 2010).

References

- AGHION, P., AND P. BOLTON (1987): “Contracts as Barriers to Entry,” *American Economic Review*, 77(3), 388–401.
- ASKER, J. (2005): “Diagnosing Foreclosure due to Exclusive Dealing,” Working Paper.
- ASKER, J., AND H. BAR-ISAAC (2014): “Raising Retailers’ Profits: On Vertical Practices and the Exclusion of Rivals,” *American Economic Review*, 104(2).
- BERNHEIM, B. D., AND M. WHINSTON (1998): “Exclusive Dealing,” *Journal of Political Economy*, 106(1), 64–103.
- BERRY, S. (1994): “Estimating discrete-choice models of product differentiation,” *RAND Journal of Economics*, 25(2), 242–261.
- BORK, R. (1978): *The Antitrust Paradox*. New York: Free Press.
- BRICKLEY, J., AND F. DARK (1987): “The Choice of Organizational Form: The Case of Franchising,” *Journal of Financial Economics*, 18, 401–420.
- CHAO, Y., AND G. TAN (2013): “All-Units Discounts, Quantity Forcing, and Capacity Constraint,” Working Paper.
- CONLON, C., AND J. H. MORTIMER (2013a): “Demand Estimation Under Incomplete Product Availability,” *American Economic Journal: Microeconomics*, 5(4), 1–30.
- CONLON, C. T., AND J. H. MORTIMER (2010): “Effects of Product Availability: Experimental Evidence,” NBER Working Paper No. 16506.
- (2013b): “An Experimental Approach to Merger Evaluation,” NBER Working Paper No. 19703.
- DENECKERE, R., H. MARVEL, AND J. PECK (1996): “Demand Uncertainty Inventories and Resale Price Maintenance,” *Quarterly Journal of Economics*, 111(3), 885–914.
- (1997): “Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition,” *American Economic Review*, 87(4), 619–641.
- ENGBER (2009): “Let Them Drink Water! What a fat tax really means for America.,” <http://www.slate.com/id/2228713/>.
- FUMAGALLI, C., AND M. MOTTA (2006): “Exclusive Dealing and Entry, when Buyers Compete,” *American Economic Review*, 96(3), 785–795.
- HOTZ, V. J., AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60(3), 497–529.

- KLEIN, B., AND K. M. MURPHY (1988): “Vertical Restraints as Contract Enforcement Mechanisms,” *Journal of Law and Economics*, 31(2), 265–297.
- KOLAY, S., G. SHAFFER, AND J. A. ORDOVER (2004): “All-Units Discounts in Retail Contracts,” *Journal of Economics and Management Strategy*, 13(3), 429–459.
- LAFONTAINE, F. (1992): “Agency Theory and Franchising: Some Empirical Results,” *RAND Journal of Economics*, 23, 263–283.
- McFADDEN, D. (1978): “Modelling the Choice of Residential Location,” in *Spatial Interaction Theory and Planning Models*, ed. by A. Karlqvist, L. Lundsqvist, F. Snickars, and J. Weibull. North-Holland.
- O’BRIEN, D. P. (2013): “All-Units Discounts and Double Moral Hazard,” Working Paper.
- POSNER, R. (1976): *Antitrust Law: An Economic Perspective*. University of Chicago Press.
- RASMUSEN, E. B., J. M. RAMSEYER, AND J. WILEY, JOHN S. (1991): “Naked Exclusion,” *The American Economic Review*, 81(5), pp. 1137–1145.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033.
- SEGAL, I. R., AND M. D. WHINSTON (2000): “Naked Exclusion: Comment,” *American Economic Review*, 90(1), 296–309.
- SHAFFER, G. (1991a): “Capturing Strategic Rent: Full-Line Forcing, Brand Discounts, Aggregate Rebates, and Maximum Resale Price Maintenance,” *Journal of Industrial Economics*, 39(5), 557–575.
- (1991b): “Slotting Allowances and Resale Price Maintenance: A Comparison of Facilitating Practices,” *RAND Journal of Economics*, 22(1), 120–135.
- SHEPARD, A. (1993): “Contractual Form, Retail Price, and Asset Characteristics in Gasoline Retailing,” *RAND Journal of Economics*, 24(1), 58–77.
- SMITH, S. (2010): “City may curb sales of sugary beverages,” *The Boston Globe*, September 20th.
- STOKEY, N. L., R. LUCAS, AND E. PRESCOTT (1989): *Recursive Methods for Economic Dynamics*. Harvard University Press.
- SUDHIR, K., AND V. R. RAO (2006): “Are Slotting Allowances Efficiency-Enhancing or Anti-Competitive?,” *Journal of Marketing Research*, 43(2), 137–155.
- TELSER, L. (1960): “Why Should Manufacturers Want Fair Trade?,” *Journal of Law and Economics*, 3, 86–105.

TIROLE, J. (1988): *The Theory of Industrial Organization*. MIT Press.

TRAIN, K. (2003): *Discrete Choice Methods with Simulation*. Cambridge University Press.




Figure 1: Conventional Quantity Discount and All Units Discount

The Only Candy You Need To Stock In Your Machine!

Spiral#1	Spiral#2	Spiral#3	Spiral#4	Spiral#5	Spiral#6	Spiral#7	Spiral#8
							
M & M 'S® Peanut Candies	SNICKERS® Bar	Twix® Cookie Bar	3 MUSKETEERS® Bar	MILKY WAY® Bar	M & M 'S® Milk Chocolate Candies	SKITTLES® Candies Original	STARBURST® Fruit Chews Original
#1 Selling Confection Item in Vending!	#2 Selling Confection Item in Vending!	#3 Selling Confection Item in Vending!	#4 Selling Confection Item in Vending!	#11 Selling Confection Item in Vending!	#6 Selling Confection Item in Vending!	#5 Selling Confection Item in Vending!	#9 Selling Confection Item in Vending!

- Based on the current business environment, vend operators are looking for one supplier to cover all of their Candy needs

- ▶ MARS - 100% Real Chocolate!
- ▶ MARS - 100% Real Sales!



Proven 52 Weeks Ending 10/4/09

MARS
chocolate
north america

2010 Vend Operator Program

Gold Rebate Level

- Continuously stock 6 Singles or King Size items

- ▶ Reduction from 7 must-stock items in 2009!

- SNICKERS® Bar singles or king size
- M&M'S® Peanut Chocolate Candies singles or king size
- M&M'S® Candies – any other variety (Milk Chocolate, Almond, Peanut Butter or Coconut) singles or king size
- TWIX® Cookie Bar single – any variety singles or king size
- 3 MUSKETEERS® Bar or MILKY WAY® Bar- any variety singles or king size
- SKITTLES® Bite Size Candies or STARBURST® Fruit Chews – any variety singles or king size

- Index >90 versus 2009

- ▶ Quarterly case index of 90 versus 2009

MARS
chocolate
north america

32

GOLD Rebate	SNICKERS® and M&M'S® Peanut	5%	\$ 0.015
Index 90	All other MARS vend items	8%	\$ 0.040

Figure 2: Mars Vend Operator Rebate Program




Figure 3: Histogram of Daily Sales for Machines in Group D




Figure 4: Profits Per Consumer as a Function of Restocking Policy

Table 1: Results from varying capacity in numerical toy example

Threshold $\bar{q} = 65$	Capacity = 45	Capacity = 65
Market:		
Total Sales([1,2]) > Total Sales([1,3])	66.72%	24.16%
Total Sales([1,2]) < Total Sales([1,3])	26.49%	66.59%
Total Sales([1,2]) = Total Sales([1,3])	6.79%	9.25%
Mean(Sales([1,2])-Sales([1,3]))	2.92	-2.00
as percent of sales	3.36%	-2.17%
Retailer:		
Retailer prefers [1,2]	25.80%	0.61%
Retailer prefers [1,3], No Rebate	0.33%	0.69%
Retailer prefers [1,3]	74.20%	99.39%
Mean Retailer profit([1,2])	60.41	63.21
Mean Retailer profit([1,3]), No Rebate	50.36	56.68
Mean Retailer profit([1,3])	62.88	70.85
Firm A:		
Firm A prefers [1,2]	0.19%	13.78%
Firm A prefers [1,3]	99.55%	85.43%
Mean Firm A profit under [1,2]	17.98	21.52
Mean Firm A profit under [1,3]	20.96	23.62
Firm A pays rebate under [1,2]	0%	0%
Firm A pays rebate under [1,3]	100%	100%
Firm B:		
Mean Firm B profits under [1,2]	8.36	7.73

*Under the assumed demand patterns, retailer always stocks product 1.

Table 2: Comparison of National Marketshares with Experimental Firm

Natl Rank	Manuf	Product	Avail	Share
1	Mars	M&M Peanut*	96	23.0
2	Mars	Snickers*	96	22.0
3	Mars	Twix Bar*	79	13.0
4	Hershey	Reeses Peanut Butter Cups	29	3.7
5	Mars	Three Musketeers*	34	4.3
6	Mars	Skittles*	77	6.5
7	Mars	M&M Milk Chocolate*	47	6.4
8	Mars	Starburst	16	1.0
9	Nestle	Butterfinger	33	2.7
10	Hershey	w/Almond	0	0
11	Mars	Milky Way/Other	33	2.6
17	Nestle	Raisinets*	78	8.9

National Shares: Mars 64%, Hershey 24%

MarkVend Shares: Mars 80%, Hershey 8.5%

MarkVend averages 6.86 confection facings per machine

Table 3: Products used in Counterfactual Analyses

M&M Peanut
M&M Milk Chocolate
Snickers
Twix Caramel
Raisinets
Strwbry Pop-Tarts
Oat n Honey Granola Bar
Grandmas Choc Chip
Choc Chip Famous Amos
Rasbry Knotts
Ritz Bits Chs Vend
Ruger Vanilla Wafer
Kar Sweet & Salty Mix
Farleys Mixed Fruit Snacks
Planters Salted Peanuts 2 oz LSS
Zoo Animal Cracker Austin
Rold Gold Pretzels LSS
Snyders Nibblers F/F
Ruffles Ched/SC LSS
Cheez-It Original SS
Frito LSS
Dorito Nacho LSS
Cheeto LSS
Smartfood LSS
Sun Chip LSS
Lays Potato Chips 1oz SS
Baked Lays LSS
Munchos Potato Chips
Hot Stuff Jays

Table 4: Group Mean and Standard Deviation

	Group Size	Vends/Visit		Revenue/Visit		Avg Sales/Day	
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

Table 5: Results from Snickers and M&M Peanut Joint Experiment

Product	Control	Treatment	Change	% Change	Difference	T-Stat	Obs
Assorted Chocolate	104.5	227.8	123.2	117.9	1.79	6.12	69
Twix Caramel	213.0	313.3	100.3	47.1	1.43	5.64	70
Reeses Peanut Butter Cups	109.0	202.2	93.3	85.6	1.23	4.30	76
Assorted Pastry	287.4	374.2	86.9	30.2	1.16	3.60	75
M&M Milk Chocolate	132.0	196.9	64.9	49.2	1.18	3.59	55
Assorted Nuts	359.3	415.8	56.6	15.7	0.73	2.28	78
Assorted Cookie	314.7	359.3	44.6	14.2	0.51	1.75	88
Assorted Nonchocolate Candy	263.4	301.1	37.7	14.3	0.45	1.80	83
Assorted Chips	548.2	585.6	37.4	6.8	0.43	1.35	87
Raisinets	184.0	215.9	31.9	17.3	0.44	1.99	73
Choc Chip Famous Amos	227.0	241.2	14.1	6.2	0.16	0.73	89
Rasbry Knotts	70.7	79.7	8.9	12.6	0.11	0.82	79
Assorted Pretzel/Popcorn	962.0	969.8	7.8	0.8	0.09	0.24	89
Assorted Fruit Snack	103.6	107.7	4.1	4.0	0.06	0.31	71
Dorito Nacho LSS	284.5	282.6	-1.9	-0.7	-0.02	-0.10	89
Baked Chips (Con)	262.8	255.8	-7.0	-2.7	-0.08	-0.35	88
Assorted Cracker	114.4	93.3	-21.1	-18.5	-0.28	-1.18	75
Sun Chip LSS	198.1	174.6	-23.5	-11.9	-0.29	-1.34	80
Cheeto LSS	349.8	325.7	-24.1	-6.9	-0.27	-1.38	89
Assorted Salty Snack	711.9	678.1	-33.9	-4.8	-0.38	-1.16	89
Assorted Energy	272.1	229.0	-43.1	-15.8	-0.61	-1.90	71
Zoo Animal Cracker Austin	292.1	235.0	-57.1	-19.6	-0.64	-3.18	89
Snickers	379.4	13.2	-366.2	-96.5	-4.11	-16.00	89
M&M Peanut	425.9	9.4	-416.5	-97.8	-4.68	-18.19	89
Total	7,170.0	6887.3	-282.7	-3.9	-3.18	-12.07	89

Table 6: Top 5 Substitutes (Vends)

Snickers	Peanut M&M	Both
M&M Peanut*	Snickers*	Assorted Chocolate*
Twix Caramel*	Assorted Pretzel/Popcorn*	Twix Caramel*
Assorted Pretzel/Popcorn	Assorted Nuts*	Reeses Peanut Butter Cups*
Assorted Nuts*	Twix Caramel*	Assorted Pastry*
Assorted Fruit Snack*	M&M Milk Chocolate*	M&M Milk Chocolate*
Focal (-535.0)	Focal (-605.5)	Focal (-782.7)
Top 5 (539.9)	Top 5 (494.3)	Top 5 (468.6)
Total (-216.8)	Total (-197.6)	Total (-282.7)

Table 7: Downstream Profit Impact

			Before Rebate			After Rebate		
Experiment	Vends	Obs	Difference In:		T-Stat	Difference In:		T-Stat
			Margin	Profit	of Diff	Margin	Profit	of Diff
Snickers	-216.82	109	0.39	-56.75	-2.87	0.24	-73.26	-4.33
Peanut M&Ms	-197.58	115	0.78	-10.74	-0.58	0.51	-39.37	-2.48
Double	-282.66	89	1.67	-4.54	-0.27	1.01	-54.87	-3.72

Table 8: Upstream (Manufacturer) Profits

Experiment	Mars	Hershey	Nestle	Other	% Before	% After
Snickers	-26.37	5.89	19.32	-20.26	31.7%	11.9%
Peanut M&Ms	-68.38	32.76	11.78	-9.36	86.4%	50.2%
Snickers + Peanut M&Ms	-130.81	61.43	20.22	37.10	96.7%	59.5%

Table 9: Parametric Model Estimates

	Random Coefficients		Nested Logit	
σ_{Salt}	0.506 [.006]	0.458 [.010]		
σ_{Sugar}	0.673 [.005]	0.645 [.012]		
σ_{Peanut}	1.263 [.037]	1.640 [.028]		
$\lambda_{Chocolate}$			0.828 [.003]	0.810 [.005]
$\lambda_{CandyNon-Choc}$			0.908 [.007]	0.909 [.009]
$\lambda_{Cookie/Pastry}$			0.845 [.004]	0.866 [.006]
λ_{Other}			0.883 [.005]	0.894 [.006]
$\lambda_{SaltySnack}$			0.720 [.003]	0.696 [.004]
# Nonlinear Params	3	3	5	5
Product FE	73	73	73	73
# Fixed Effects ξ_t	15256	2710	15256	2710
Total Parameters	15332	2786	15334	2788
LL	-4372750	-4411184	-4372147	-4410649
Total Sales	2960315	2960315	2960315	2960315
BIC	8973960	8863881	8972783	8862840
AIC	8776165	8827939	8774962	8826873

Table 10: Stocking Policies (Restock after how many expected sales?)

Additional Products Stocked*:	Retailer	Rebate	Integrated	Industry	Retail Understock	Rebate Increase
Reeses(H), 3 Musketeers(M)	267	262	246	249	8.5%	2.0%
Reeses(H), Payday(H)	263	257	239	243	10.0%	2.5%
3 Musketeers (M), Milkyway(M)	264	260	244	246	8.2%	1.6%
Reeses(H), Butterfinger(N)	263	258	240	244	9.6%	2.1%
Butterfinger(N), Crunch(N)	257	252	233	236	10.3%	2.1%

*The five base products are: Snickers, Peanut M&M, Twix, Plain M&M (Mars), and Raisinets (Nestle). Manufacturers are denoted as (H) for Hershey's, (N) for Nestle, and (M) for Mars.

Table 11: Profits under Alternate Product Assortments and Stocking Policies

Policy		Retail (No Reb.)	Rebate	Mars	Hershey	Nestle	Integrated	Industry	Consumer
Reeses PB Cup(H), 3 Musketeers(M)									
Retailer	267	36,398		11,763	1,305	1,260	48,161	50,726	63,371
Rebate	263	36,395		11,789	1,302	1,258	48,184	50,744	63,425
Integrated	246	36,342		11,885	1,293	1,249	48,227	50,769	63,576
Industry	249	36,356		11,870	1,294	1,251	48,226	50,771	63,559
Reeses PB Cup(H), Payday(H)									
Retailer	263	36,668		10,091	2,173	1,284	46,759	50,216	62,600
Rebate	258	36,665		10,128	2,168	1,282	46,793	50,242	62,663
Integrated	239	36,596		10,253	2,151	1,272	46,849	50,272	62,801
Industry	243	36,618		10,229	2,154	1,274	46,848	50,276	62,787
3 Musketeers(M), Milkyway(M)									
Retailer	265	36,101	2,096	13,100	0	1,259	49,201	50,460	63,038
Rebate	261	36,099	2,100	13,123	0	1,257	49,222	50,479	63,092
Integrated	245	36,052	2,113	13,208	0	1,249	49,260	50,509	63,236
Industry	246	36,057	2,112	13,203	0	1,250	49,260	50,510	63,231

Table 12: Profits after Mars-Hershey Merger

Policy		Retail (No Reb.)	Rebate	Mars/ Hershey	Nestle	Integrated	Industry	Consumer
Reeses PB Cup(H), 3 Musketeers (M)								
Retailer	267	36,063	2,144	13,403	1,260	49,466	50,726	63,371
Rebate	263	36,060	2,149	13,426	1,258	49,486	50,744	63,425
Integrated	248	36,020	2,160	13,501	1,250	49,521	50,771	63,565
Industry	249	36,024	2,159	13,497	1,251	49,521	50,771	63,559
Butterfinger (N), Crunch (N)								
Retailer	257	36,601		10,228	2,705	46,830	49,535	61,683
Rebate	252	36,597		10,267	2,700	46,864	49,564	61,743
Integrated	234	36,532		10,389	2,682	46,921	49,602	61,857
Industry	236	36,543		10,377	2,684	46,920	49,603	61,853

Table 13: Profits after Mars-Nestle Merger

Policy		Retail (No Reb.)	Rebate	Mars/ Nestle	Hershey	Integrated	Industry	Consumer
Reeses PB Cup(H), 3 Musketeers(M)								
Retail	267	36,119		13,302	1,305	49,421	50,726	63,371
Rebate	263	36,117		13,325	1,302	49,442	50,744	63,425
Integrated	247	36,071		13,406	1,293	49,477	50,770	63,571
Industry	249	36,080		13,397	1,294	49,477	50,771	63,559
Reeses PB Cup(H), Payday(H)								
Retailer	263	36,384		11,659	2,173	48,044	50,216	62,600
Rebate	258	36,381		11,693	2,168	48,074	50,242	62,663
Integrated	241	36,326		11,796	2,152	48,122	50,274	62,795
Industry	243	36,336		11,785	2,154	48,121	50,276	62,787
3 Musketeers(M), Milkyway(M)								
Retailer	264	35,822	2,343	14,643	0	50,465	50,465	63,052
Rebate	261	35,821	2,345	14,658	0	50,479	50,479	63,092
Integrated	246	35,781	2,356	14,729	0	50,510	50,510	63,231
Industry	246	35,781	2,356	14,729	0	50,510	50,510	63,231

Table 14: Profits after Hershey-Nestle Merger

Policy	Retail	Rebate	Mars	Hershey/ Nestle	Integrated	Industry	Consumer
	(No Reb.)						
	Reeses PB Cup(H), 3 Musketeers						
Retail	267	36,398	11,763	2,565	48,161	50,726	63,371
Rebate	263	36,395	11,789	2,560	48,184	50,744	63,425
Integrated	246	36,342	11,885	2,542	48,227	50,769	63,576
Industry	249	36,356	11,870	2,545	48,226	50,771	63,559
	Reeses PB Cup(H), Payday(H)						
Retail	263	36,668	10,091	3,457	46,759	50,216	62,600
Rebate	258	36,665	10,128	3,450	46,793	50,242	62,663
Integrated	239	36,596	10,253	3,422	46,849	50,272	62,801
Industry	243	36,618	10,229	3,428	46,848	50,276	62,787
	3 Musketeers(M), Milkyway(M)						
Retail	265	36,101	2,096	13,100	1,259	49,201	63,038
Rebate	261	36,099	2,100	13,123	1,257	49,222	63,092
Integrated	245	36,052	2,113	13,208	1,249	49,260	63,236
Industry	246	36,057	2,112	13,203	1,250	49,260	63,231

Table 15: Comparison under Alternate Ownership Structures

Alternative	(H,H)	(H,N)	Mars-Hershey (N,N)	Mars-Nestle (H,H)	Hershey-Nestle (H,H)	(H,N)
Retail Profit	36,668	36,649	36,601	36,384	36,668	36,649
Retail Profit (post Rebate)	38,165	38,165	38,180	38,137	38,165	38,165
Δ Retail	1,497	1,516	1,579	1,753	1,497	1,516
Δ Competitor	-2,173	-1,316	-1,455	-2,173	-2,208	-2,262
Combined Change	-676	200	124	-420	-711	-746
Price to Avoid Exclusion	13.3	none	none	8.3	24.2	24.3
Possible Rebate Reduction	none	9.5%	4.6%	none	none	none
Δ Mars	1,004	1,021	1,113	714	1,004	1,021
Δ Industry	293	275	1,236	294	293	275
Δ CS	636	602	1,882	631	636	602
Δ CS (Efficiency)	198	198	194	179	198	198
Δ Ind Max	-260	-260	0	-260	-260	-260
Δ CS Max	-340	-340	0	-340	-340	-340

Notes: Table compares the welfare impacts of an exclusive Mars stocking policy with the alternatives identified in tables 11-14.

Appendix

A.1: Solving for Optimal Assortment

Even the simple problem faced by the retailer in (1) can be extremely challenging to solve. Begin by profiling out the optimal effort level conditional on $e(a)$, the retailer chooses a vector a with elements a_j to solve:

$$\max_a \pi_r(a, e(a)) - c(e(a)) \text{ s.t. } \sum_j a_j \leq k \quad \text{with } a_j \in \{0, 1\}$$

This is a binary integer programming problem, and it is quite challenging since there are $\binom{J}{k}$ possible product assortments to consider. Even in the case where $\pi(a) - c(a)$ is a linear function of the vector a , the problem is NP-hard. Once we allow for substitutes or complements in the profit function, only brute-force solutions exist. This means for each possible a , we must solve the dynamic programming problem and calculate profits. For the case of $\binom{12}{7}$, this represents 792 possible combinations that must be considered.

To avoid considering all possible combinations, we consider several heuristics in our empirical work. One useful heuristic is the “greedy heuristic” where the retailer adds the next most profitable product in each slot. In order of profitability of Mars products, the ranking is (M&M Peanut, Snickers, Twix, M&M Plain, 3 Musketeers, Milkyway, Skittles). We do not present results that consider removing M&M Peanut, Snickers, Twix, M&M Plain, or Raisinets from the choice set. In practice, we consider robustness to replacing M&M Plain or Raisinets, though we do not present those results because they are always strictly worse than the ones we do report. This implies the retailer’s decision is over only $\binom{7}{2} = 21$ alternatives which is much more manageable, again we reported only the best alternatives among those 21.

A.2: Solving for Endogenous Linear Prices

Allowing for endogenous linear prices (even in the absence of rebates) is a challenging extension. In most vertical settings, the retailer stocks all products and chooses retail prices given a set of wholesale prices. In our setting, the retail price is fixed, and the retailer chooses a product assortment by solving the binary integer programming problem described above. Even with three wholesale prices that are the same across all of a manufacturer’s products (w_m, w_h, w_n), this is a challenging problem, because the heuristic solutions employed in the previous section no longer work.

This problem arises because small changes in one of the wholesale prices w_m for example can result in discrete changes in the product assortment, and discontinuous jumps in the upstream profits. For Bertrand competition in homogeneous products this results in the well-known limit pricing result. With multiple differentiated products, this implies that First-Order Conditions are not sufficient to characterize the equilibrium, and that many such equilibria may exist.

A.3: Computing Treatment Effects

One goal of the product removals is to determine how product-level sales respond to changes in availability. Recall from the text that q_{jt} denotes the sales of product j in machine-week t , superscript 1 denotes sales when a focal product(s) is removed, superscript 0 denotes sales when a focal product(s) is available, the set of available products is A , and F is 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 a treatment week, and control week respectively, and $q_{fs}^0 = \sum_{j \in F} q_{js}^0$ is the sales of the removed products. 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. This can be seen in Figure 6, which plots the overall sales of all machines in our sample on a weekly basis. 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 experiments were run 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 stock-out events, 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:

$$\{s : s \neq t, Q_t^0 - Q_s^1 \in [0, q_{fs}^0]\} \quad (9)$$

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 product removal, which would be problematic for our study.

We propose a slight modification of (9) which removes the bias. That is, we can replace q_{fs}^0 with $\widehat{q_{fs}^0} = E[q_{fs}^0 | Q_s^0]$. An easy way to obtain the expectation is to run an OLS regression of q_{fs}^0 on Q_s^0 , at the machine level and use the predicted value. This has the nice property that the error is orthogonal to Q_s^0 , which ensures that our choice of weeks is now 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:

$$S_t = \{s : s \neq t, A_t \approx A_s, Q_t^0 - Q_s^1 \in [0, \hat{b}_0 + \hat{b}_1 Q_s^0]\} \quad (10)$$

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

$$\Delta q_{jt} = q_{jt}^1 - \frac{1}{\#S_t} \sum_{s \in S_t} q_{js}^0 \quad (11)$$

While this approach has the advantage that it generates substitution patterns consistent with consumer theory, it may be 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.⁴²

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 (11) to construct a set of pseudo-observations for the difference, and employ a paired t-test.

⁴²Weeks in which the other five treatments were run (for the salty-snack and cookie categories) are excluded from the set of potential control weeks.




Figure 5: Overall Weekly Sales at Site 93