

Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence ^{*}

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

(Click Here for Current Version)

April 20, 2105

Abstract

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

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

[†]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

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

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

We address these challenges by examining a vertical rebate known as an All-Units Discount (AUD). The specific AUD we study is used by the dominant chocolate candy man-

¹Depending on the form of the vertical rebate, it may be referred to as a loyalty contract (if it uses a market-share requirement) or an all-units discount (if it is based on a retailer's total purchases). The use of a vertical rebate contract was central to several recent antitrust cases involving Intel. 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 2014, *Eisai v. Sanofi-Aventis* applied the *Meritor* reasoning to loyalty contracts between Sanofi and hospitals for the purchase of a blood-clotting drug, ruling in favor of the drug manufacturer on the basis of a predatory pricing standard. The DOJ and the FTC, in June 2014, held a joint workshop on Conditional Pricing Practices exploring the implications of this class of contracts for antitrust policy.

ufacturer in the United States: Mars, Inc.² The AUD implemented by Mars consists of three main features: a retailer-specific per-unit discount, a retailer-specific quantity target or threshold, and a ‘facing’ requirement that the retailer carry at least six Mars products. Mars’ AUD stipulates that if a retailer meets the facing requirement and his total purchases exceed the quantity target, Mars pays the retailer an amount that is equal to the per-unit discount multiplied by the retailer’s total quantity purchased. We examine the effect of the rebate contract through the lens of a retail vending operator, Mark Vend Company, for whom we are able to collect extremely detailed information on sales, wholesale costs, and contractual terms. The retailer also agreed to run a large-scale field experiment on our behalf, which provides us with additional insight into how the AUD might influence the retailer’s decisions. To the best of our knowledge, no previous study has had the benefit of examining a vertical rebate contract using such rich data and exogenous variation.

The insights that we gain from studying Mars’ rebate contract allow us to contribute to understanding principle-agent models in which downstream moral hazard plays an important role. Downstream moral hazard is an important feature of many vertically-separated markets, and is thought to drive a variety of vertical arrangements such as franchising and resale price maintenance (RPM).³ However, empirically measuring the effects of downstream moral hazard is difficult. Downstream effort may be impossible to measure directly, and vertical arrangements are endogenously determined, making it difficult to identify the effects of downstream moral hazard on upstream firms. Our ability to exogenously vary the result of downstream effort (in this case, retail product availability), combined with detailed data on wholesale prices, allows us to directly document the effects of downstream moral hazard on the revenues of upstream firms.

In order to analyze the effect of Mars’ AUD contract, we specify a 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.⁴ The number of units the retailer can stock for each product is constrained by the capacity of his vending

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

³See, among others, Shepard (1993) for an early empirical study of principle-agent problems in the context of gasoline retailing, and Hubbard (1998) for an empirical study of a consumer-facing principle-agent problem.

⁴By 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.

machines, and we interpret retailer effort as the frequency with which the retailer restocks his machines. In order to calculate the retailer’s optimal effort level, we compute a dynamic restocking model à la Rust (1987), in which the retailer chooses how long to wait between restocking visits.⁵ Due to the capacity constraints of a vending machine, the number of unique products the retailer can stock is relatively small. Thus, we compute the dynamic restocking model for several discrete sets of products, and we assume that the retailer chooses to stock the set of products that maximizes his profits. These features of the 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.⁶

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 outcome 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, increases the retailer’s share of the cost of stock-out events from around 10% to nearly 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 effort and assortment decisions. Mars’ AUD contract is designed to induce greater retailer effort through more frequent re-stocking.⁷ However, when the retailer increases his re-stocking effort under the contract, he re-stocks all products, regardless of manufacturer. Over the relevant range of the retailer’s re-stocking policy, more frequent re-stocking *reduces* sales of Hershey and Nestle products, because these products

⁵Rather than assuming retailer wait times are optimal and using the dynamic model to estimate the cost of re-stocking, we do the reverse: we use an outside estimate of the cost of re-stocking based on wage data from the vending operator, and use the model to compute the optimal wait time until the next restocking visit.

⁶These features also characterize other industries, such as brick-and-mortar retail and live entertainment.

⁷Downstream effort in other settings might also include promotional activities or price reductions.

no longer benefit from forced substitution when the dominant Mars products sell out. Thus, downstream effort is substitutable across upstream firms and implies that Hershey and Nestle do not have an incentive to offer a rebate of the same form.⁸ We provide evidence that the rebate induces greater retailer effort, and that the level of retailer effort that is optimal from the perspective of a hypothetical vertically-integrated Mars-retailer entity is slightly higher than the effort level preferred by the industry as a whole.

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 in the absence of the AUD contract. The ability to do this is an indication of possible foreclosure. We find evidence that the retailer can increase profits by substituting a Hershey product for a Mars product, but that the threat of losing the rebate discourages him from doing so. Thus, the AUD does result in foreclosure.

In spite of the evidence of foreclosure and the failure of the AUD to implement the industry-optimal level of downstream effort, the overall welfare effects of the AUD depend on what Mars would do in the absence of the AUD. We consider three counterfactual scenarios. First, we measure the effect of dropping the AUD, holding wholesale prices fixed. Under this scenario, the retailer drops some Mars products and reduces his restocking effort. As a result, Mars and the retailer are both worse off, and consumer surplus falls. Hershey and Nestle both benefit from the reduced retailer effort and alternative product assortment, but these effects are small, and overall, social welfare goes down when the AUD is removed. Second, we drop the AUD but allow Mars to re-optimize its wholesale price against the wholesale prices of its competitors, which we hold fixed. In this scenario, Mars is worse off (compared to the current outcome with the AUD), but the retailer benefits. The change to consumer surplus is quite small, and whether the change is positive or negative depends on the level of the quantity threshold under the AUD.

In our third and final exercise, 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. Thus, we simulate the impact of various upstream mergers on the willingness of the dominant firm to offer rebate contracts, and the impact that this has on social welfare. Interestingly, we find conditions under which an upstream merger of a dominant firm (Mars) with a close competitor (Hershey) can lead to socially-efficient downstream effort and product assortment. This happens because the merger addresses the demand externalities that lead to

⁸If the retailer reduces his effort below this range, all products stock out, so that more frequent restocking increases sales of all products, including those of Hershey and Nestle. Downstream effort would be complementary in this range, but it is not profitable for the retailer.

the substitutability of retail effort across Mars and Hershey products. We also find that an upstream merger of two smaller rivals (Hershey and Nestle) can bid up the price of a downstream firm’s shelf space, even though it cannot necessarily prevent exclusion.

Having estimated the demand and supply-side models, one can, in principle, conduct a wider range of counterfactual exercises. For example, one can examine alternative contracts, such as two-part tariffs or conventional quantity discounts. Rather than pursue a list of these alternative contractual forms, we consider whether or not there is any action that Hershey could take to avoid exclusion when facing the Mars AUD contract. We find that the answer is, in many cases, “no.”

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).⁹ An important theoretical development on the potential foreclosure effects of vertical contracts is the so-called *Chicago Critique* of Bork (1978) and Posner (1976), which makes the point that because the downstream firm must be compensated for any exclusive arrangement, one should only observe exclusion in cases for which it maximizes industry profits. Much of the subsequent theoretical literature focuses on demonstrating that the *Chicago Critique’s* predictions are a bit special.¹⁰ A separate, but related, 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).¹¹ A broader

⁹In addition, 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. For situations in which retailers have the ability to set prices, Klein and Murphy (1988) show that without vertical restraints, retailers “will have the incentive to use their promotional efforts to switch marginal customers to relatively known brands...which possess higher retail margins.”

¹⁰For 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 exclusion for which industry profits are not maximized; while Bernheim and Whinston (1998) show that the *Chicago Critique* ignores externalities across buyers, and that once externalities are accounted for, it is again possible to generate exclusion that fails to maximize industry profits. 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).

¹¹This 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.

literature has also examined the conditions under which bilateral contracting might lead to (perhaps partial) exclusion.¹²

Recent theoretical work related to AUDs specifically includes Kolay, Shaffer, and Ordovery (2004), which shows 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.¹³ More recently, Chao and Tan (2014) show that AUD and quantity-forcing contracts can be used to exclude 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.

We depart from the basic theoretical framework of the *Chicago Critique* of Bork (1978) and Posner (1976) in some key ways. First, we allow for downstream moral hazard and potential efficiency gains, similar to much of the later theoretical work on vertical arrangements. Second, we study an environment in which the degree of competition across upstream firms may vary across the potential sets of products carried by the retailer, because upstream firms own multiple, differentiated products. Finally, we restrict the retailer to carrying a fixed number of these differentiated products.¹⁴

The theoretical literature following the *Chicago Critique* focuses on a wide range of settings when considering the potential effects of vertical contracts. Specifically, this literature has studied contracts used by dominant vs. non-dominant firms, contracts that do or do not reference rivals, contracts for which downstream price competition is a major concern for upstream firms (or not), and contracts that apply to single products vs. multiple products. Our setting provides empirical evidence on a vertical rebate used by a dominant firm covering multiple products, for which excessive downstream price competition is not a concern. Although the contract does not explicitly reference rivals, the facing requirement, combined with the typical capacity constraints of most vending machines, effectively limits the presence of competing brands.

One challenge for understanding the effects of vertical arrangements across this wide range of settings is that empirical evidence has primarily been available only through the course of litigation. This has the potential effect that debates about these contracts may be based on a selected sample. An important distinction of our setting is that we study a

¹²Some key examples include Rasmusen, Ramseyer, and Wiley (1991), Segal and Whinston (2000), and more recently Asker and Bar-Isaac (2014) and Chen and Shaffer (2014).

¹³In addition, Elhauge and Wickelgren (2012) and Elhauge and Wickelgren (2014) explore the potential of loyalty contracts to soften price competition, and Figueroa, Ide, and Montero (2014) examines the role that rebates can play as a barrier to inefficient entry.

¹⁴This contrasts with the “naked exclusion” of Rasmusen, Ramseyer, and Wiley (1991), in which there is a single good.

contract that has not been litigated, and for which we have detailed information on contract terms and exogenous variation in the results of the retailer’s effort. Although the welfare effects of vertical rebate contracts in other situations may differ from the impacts we estimate in our setting, we hope that our work provides a road-map for how to model the impacts of these contracts empirically.

Outside of the theoretical literature on vertical rebate contracts, our work also connects to the empirical literature on the impacts of vertical arrangements. One strand of this literature examines issues of downstream moral hazard in the context of vertical integration and the boundaries of the firm, rather than through vertical contracts per se.¹⁵ More recently, another strand of this literature examines exclusive contracts, without necessarily focusing on downstream moral hazard or effort decisions.¹⁶ The most closely-related empirical work is work on vertical bundling in the movie industry, and on vertical integration in the cable television industry. The case of vertical bundling, known as full-line forcing, is studied by Ho, Ho, and Mortimer (2012a) and Ho, Ho, and Mortimer (2012b), which examine the decisions of upstream firms to offer bundles to downstream retailers, the decisions of retailers to accept these ‘full-line forces,’ and the welfare effects induced by the accepted contracts. The case of vertical integration is studied by Crawford, Lee, Whinston, and Yurukoglu (2015), which examines efficiency and foreclosure effects of vertical integration between regional sports networks and cable distributors. A distinction between our work and Crawford, Lee, Whinston, and Yurukoglu (2015) is that we examine the potential for upstream foreclosure (i.e., manufacturers being denied access to retail distribution), while that study examines the potential for downstream foreclosure (i.e., distributors not having access to inputs).¹⁷

The rest of the paper proceeds as follows. Section 2 provides the theoretical framework for the model of retail behavior. Section 3 describes the vending industry, data, and the

¹⁵A few key examples that address downstream (and in some cases upstream) issues of moral hazard include Lafontaine (1992) and Brickley and Dark (1987), which study franchise arrangements, and Baker and Hubbard (2003) and Gil (2007), which study trucking and movies respectively; many other contributions are reviewed in Lafontaine and Slade (2007).

¹⁶Examples of this literature include Asker (2005), Sass (2005), and Chen (2014), which each examine the efficiency and foreclosure effects of exclusive dealing in the beer industry, and Chipty (2001) and Sinkinson (2014), which study the cable television and mobile phone markets respectively. Lee (2013) focuses on the interaction of exclusive contracts and network effects and competition between downstream firms. Lafontaine and Slade (2008) surveys this literature.

¹⁷From a methodological perspective, Crawford, Lee, Whinston, and Yurukoglu (2015) differ from us in their use of a bargaining model to describe the equilibrium carriage decisions of cable channels and downstream distributors. These carriage decisions are equivalent to a retailer’s choice of product assortment. The exogenous variation in product assortment induced by our field experiment makes a bargaining model unnecessary in our context.

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

2 Theoretical Framework

In a conventional nonlinear discount contract, the retailer pays a linear price w for the first \bar{q}_M 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 to all additional units, so that a retailer's variable cost of inputs from a rebating manufacturer is $C(q) = wq - \mathbf{1}[q > \bar{q}] \cdot \Delta \cdot q$. Figure 1 provides an illustration of the pricing schedule associated with a conventional quantity discount for a single product with a wholesale price of 1 and a discounted price of 0.75 for quantities greater than six, and an All-Units Discount with a wholesale price of 1, a threshold of 6, and a discounted price of 0.67. The structure of the AUD implies that for some quantity range, the retailer can make a lower total payment but receive more total units of the good. This use of a negative marginal cost has led some to believe that the use of an AUD is de facto evidence of anticompetitive behavior.¹⁸

A defense of AUD contracts is that they have the potential to be efficiency enhancing if the retailer is encouraged to exert costly effort required to sell the good.¹⁹ 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}_M , which triggers the transfer payment from the manufacturer to the retailer. Much like a two-part tariff, an appropriate choice of \bar{q}_M can induce an efficient level of retail effort.²⁰

Consider a simple framework with a single downstream retailer R , a dominant upstream firm M , and an upstream competitor N .²¹ The two upstream firms each sell several competing differentiated products, and set a single linear wholesale price per unit for all of their products, (w_M, w_N) .²² Taking wholesale prices as given, the dominant firm M proposes a nonlinear rebate contract, which consists of a discount and a threshold, (Δ, \bar{q}_M) , for which

¹⁸For discussion of alternative legal standards for adjudicating loyalty contracts, see Klein and Lerner (2014).

¹⁹This defense was employed by Intel in its recent antitrust cases, for example.

²⁰Related to the potential quantity-forcing effect of the threshold, lower retail prices are a non-contractible form of effort that is costly for the retailer to provide, and demand enhancing for an upstream firm.

²¹In our empirical application, we allow for two upstream competitors, but one is sufficient for conveying the intuition of our model.

²²Although demand may be different for different products sold by the same manufacturer within a product category, uniform wholesale pricing is a common feature of many markets. For example, manufacturers of many consumer packaged goods do not generally charge different prices for different products or flavors, despite different levels of demand and cost (i.e., snack foods, yogurt, and juice/beverages).

the threshold \overline{q}_M refers to total sales across all of M 's products.²³ After observing the wholesale prices and the terms of the rebate 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.²⁴ Finally, sales are realized, $\mathbf{q}(a, e)$, which depend on both the product assortment and the retailer effort level. We assume that the retailer sets 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 assortment, rather than retail prices.²⁵

We consider a single scalar version of non-contractible retailer 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:

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

where $\pi_R(a, e)$ is the variable profit of the retailer, given by $pq_M - [w_M q_M - \mathbf{1}[q_M > \overline{q}_M] \cdot \Delta \cdot q_M] + (p - w_N)q_N$, where p is the retail price common to all products, q_M and q_N are the total quantities sold for products of manufacturers M and N respectively, and $c(e)$ is the cost of retail effort. When the profit of M is increasing in effort, there is an incentive for M to offer a contract to the retailer that enhances his effort. We do not make any restrictions as to whether the profit of the upstream competitor is increasing or decreasing in retailer effort. Specifically, we consider the possibility that increased retailer effort may benefit competing upstream firms by leading to more frequent restocking of their products too, or may harm upstream competitors by preventing stock-outs of the dominant firm's brands, which might otherwise lead consumers to substitute to the competitor's products.

²³We abstract away from any potential facing requirement in this example. A facing requirement can substitute for a higher quantity threshold, but may be more difficult for the manufacturer to enforce.

²⁴For example, the number of unique products is often determined by shelf-space constraints at the retail location. For vending operators, there is a fixed number of "columns" (or coils) that are sized for candy bars. The only flexibility a vending operator has for changing the number of products in a machine arises if he stocks the same product in two columns.

²⁵For vending, uniform pricing is reinforced by technological constraints on providing change (e.g., nickels are thick, so prices requiring nickels to make change for \$1.00 are usually avoided). Other prominent examples of retail settings with fixed, uniform pricing include the theatrical and streaming markets for movies, digital download markets such as iTunes, and many consumer packaged goods. Retailers in these markets generally do not carry all possible products, so retailer assortment decisions are an important aspect of competition.

These demand externalities, which arise from changes in 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 the perspective of the industry as a whole. Whether the vertically-integrated $M - R$ firm would choose a level of effort that is too high or too low depends on whether retailer effort is a substitute or a complement to the profits of the upstream competitor. In other words, the upstream competitor might be able to “free-ride” on the enhanced effort that M ’s contract induces, or enhanced effort may lower N ’s sales.

Having specified the choice of effort, one can examine the retailer’s choice of product assortment a . The rebate contract may induce the retailer to stock more products by M and fewer products by N , or to induce the retailer to select products made by N that do not compete closely with M ’s products. The retailer can compute the optimal effort 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 $a' \neq a$.²⁶

Once we have characterized the retailer’s choice of (a, e) for a given set of wholesale prices, one 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 N is individually rational (IR) and incentive compatible (IC) for R . Conditional on a contract (Δ, \bar{q}) , one may also examine whether or not 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 N can avoid exclusion, one may also consider the amount by which M may be able to reduce the discount Δ and still obtain the same product assortment a (i.e., not violate the IR constraints of the retailer).

Using backward induction to solve for the optimal assortment, we may also consider the decision of the manufacturer to offer an AUD contract at existing prices.²⁷ We do not fully endogenize the initial wholesale prices (w_M, w_N) , because allowing w_M to freely adjust results in a continuum of equilibria in the game.²⁸ We cannot derive analytic predictions, because

²⁶For a discussion of the challenges involved in solving for optimal assortment, and a numerical example of assortment choice, holding prices fixed, please see appendix A.1.

²⁷This parallels recent work by Asker and Bar-Isaac (2014) and 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. Asker and Bar-Isaac (2014) provide a theoretical examination of practices by which upstream firms transfer profits to retailers. Their work employs Markov Perfect Equilibria using information on observed profits plus some uncertainty. Our results use information on expected profits, instead of observed profits plus uncertainty. This makes it easier to compute results and compare alternative contractual forms.

²⁸To illustrate, consider increasing the wholesale price to $(w_M + \epsilon)$ and the rebate to $(\Delta + \epsilon)$. This results

the optimal assortment $a(w_M, w_N)$ and the effort level $e(a, w_M, w_N)$ need not be smooth functions of prices. Small changes in wholesale prices can result in replacing products from one manufacturer with those of another.²⁹

The goal of the empirical section will be to measure the key quantities described in the framework above: the substitutability of products in the retail market, how the benefits of increased effort are distributed among the retail and manufacturer tiers, and whether effort serves as a substitute or complement in the profits of upstream firms.

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.³⁰ Under the 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. The sales goal for an operator is typically set for the combined sales of a manufacturer’s products, rather than for individual products. Mars’ rebate contract also stipulates a minimum number of product ‘facings’ that must be present 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 to observe the specific terms of the Mars rebate program for the vending operator with whom we worked. We include some promotional materials in figure 2.³¹ The

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}_M . If one continues increasing both the wholesale price w_M and the rebate Δ , this approaches a quantity-forcing contract with a linear tariff for any quantity in excess of \bar{q}_M in the limit. 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. 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. Downward deviations, in which M sets the wholesale price to $(w_M - \epsilon)$ and the rebate to $(\Delta - \epsilon)$, undercut the AUD’s ability to leverage previous sales to induce greater downstream effort. In the limit, this approaches simple linear pricing. Section A.2 in the appendix provides further detail on comparisons to other contracts.

²⁹Section A.3 in the appendix provides further discussion of endogenous linear prices and an illustration of this effect using our estimated model.

³⁰For confections products, Mars is the dominant manufacturer in vending, and is the only manufacturer to offer a true AUD contract. Hershey and Nestle offer wholesale ‘discounts,’ but these have a quantity threshold of zero (i.e., their wholesale pricing is equivalent to linear pricing). We do not examine beverage sales, because many beverage machines at the locations we observe are serviced directly by Coke or Pepsi.

³¹A full slide deck, titled ‘2010 Vend Program’ materials, dated December 21, 2009, is available at <http://vistar.com/KansasCity/Documents/Mars%202010%20Operatopr%20rebate%20program.pdf>. (Last

program employs the slogan *The Only Candy You Need to Stock in Your Machine!*, and specifies a facing requirement of six products and a quarterly sales target. We also observe, but cannot report, the amount of the rebate received by our retailer, Mark Vend Company. Unlike the Intel rebate program, these rebates do not explicitly condition on market share or the sales of competitors. However, most vending machines typically carry between six and eight candy bar varieties, so the facing requirement may limit shelf space for competing brands.³²

In table 1 we report the national sales ranks, availability, and shares in the vending industry for the 10 top-ranked products nationally, as well as the availability and shares for the same products from Mark Vend. There are some patterns that emerge. The first is that Mark Vend stocks some of the most popular products sold by Mars (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, even though nationally Reese’s Peanut Butter Cups is the fourth most popular product. On the other hand, Raisinets, a Nestle product, is stocked in 78% of machine weeks for Mark Vend and constitutes almost 9% of sales, despite a national rank below the top 45.

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.³³ The alternative explanation is that the rebate may induce Mark Vend to substitute from Nestle/Hershey brands to Mars brands when making stocking decisions, or that when Mark Vend does stock products from competing manufacturers (e.g., Nestle Raisinets), he chooses brands that do not steal business from key Mars products.

3.2 Data Description and Experimental Design

All of our price and quantity data are provided by Mark Vend. 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,

accessed on April 19, 2015.)

³²While there is some ability for a vending operator to adjust the overall number of candy bars in a machine, it is often difficult to do without upgrading capital equipment, because candy bars and salty snacks do not use the same size ‘slots.’

³³For example, Skittles, a fruit flavored candy sold by Mars, is primarily marketed to younger consumers.

including salty snacks, cookies, and other products in addition to confections. We observe retail and wholesale prices for each product at each service visit during our 38-month panel. There is relatively little price variation over time within a site, and almost no price variation within a product category (e.g., confections) at a site. ‘Natural,’ or non-experimental, stock-outs are extremely rare for our set of machines.³⁴ 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.³⁵

In addition to sharing the terms of his rebate contact with us, the owner of Mark Vend implemented a field experiment for us in which his drivers exogenously removed either one or two top-selling Mars confection products from a set of 66 vending machines. The product removals are recorded during each service visit to individual vending machines.³⁶ Implementation of each product removal was fairly straightforward; we removed either one or both of the two top-selling Mars products from all machines for a period of roughly 2.5 to 3 weeks. The focal products were Snickers and Peanut M&Ms.³⁷ 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 38 months, from January of 2006 until February of 2009. During each 2-3 week experimental period, most machines receive about

³⁴Mark Vend commits to a low level of stock-out events in its service contracts. This implies that much of the variation in product assortment comes either from product rotations, or our own exogenous product removals. Product rotations tend to occur primarily for ‘marginal’ products, so in the absence of exogenous variation in availability, the substitution patterns between marginal products is often much better identified than substitution patterns between continually-stocked best-selling products.

³⁵For example, we combine Milky Way Midnight with Milky Way, and Ruffles Original with Ruffles Sour Cream and Cheddar. In addition to the data from Mark Vend, we also collect data on product characteristics online and through industry trade sources. For each product, we note its manufacturer, as well as the following set of product characteristics: package size, number of servings, and nutritional information. Nutritional information includes weight, calories, fat calories, sodium, fiber, sugars, protein, carbohydrates, and cholesterol. For 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.

³⁶The machines are located in office buildings, and have substitution patterns that are very stable over time. In addition to the three treatments described here, we also ran 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.

³⁷Whenever 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.

three service visits. However, the length of service visits varies across machines, with some machines visited more frequently than others. Machines are serviced on different schedules, and as a result, it is convenient to organize observations by machine-week, rather than by visit when analyzing the results of the experiment. When we do this, we assume that sales are distributed uniformly among the business days in a service interval, and assign those business days to weeks. Different experimental treatments start on different days of the week, and we allow our definition of when weeks start and end to depend on the client site and experiment.³⁸

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

In principle, calculating the effect of product removals is straightforward. In practice, however, there are two challenges in implementing the removals and interpreting the data generated by them. First, there is 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 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.4 of the appendix.

3.3 Results of Product Removals

Our first exogenous product removal eliminated Snickers from all 66 vending machines involved in the experiment; the second removal eliminated Peanut M&Ms, and the third eliminated both products.³⁹ These products correspond to the top two sellers in the confections

³⁸For example, at some site-experiment pairs, we define weeks as Tuesday to Monday, while for others we use Thursday to Wednesday.

³⁹As noted in table 1, both Snickers and Peanut M&Ms are owned by Mars.

category, both at Mark Vend and nationwide.

We summarize substitution to the top five substitutes for all three removals in table 2.⁴⁰ In general, the substitution patterns we recover are reasonable; the top substitutes generally include Snickers or Peanut M&Ms if one of the two products is available. Twix, the third-best selling Mars brand both nationally and in our sample, is also a top substitute.⁴¹ Consumers also substitute to products outside the confections category, such as Planters Peanuts or Rold Gold Pretzels.

One of the results of the product removal is that many consumers purchase another product in the vending machine. While many of the alternative brands are owned by Mars, several of them are not. If those other brands have similar (or higher) margins for Mark Vend, substitution may cause the cost of each product removal to be distributed unevenly across the supply chain. Table 3 summarizes the impact of the product removals for Mark Vend. In the absence of any rebate payments, we see the following results. Total vends decrease by 217 units and retailer profits decline by \$56.75 when Snickers is removed. When Peanut M&Ms is removed, vends go down by 198 units, but Mark Vend's average margin on all items sold in the machine rises by 0.78 cents, and retailer revenue declines only by \$10.74 (a statistically insignificant decline). Similarly, in the joint product removal, overall vends decline by roughly 283 units, but Mark Vend's average margin rises by 1.67 cents per unit, so that revenue declines by only \$4.54 (again statistically insignificant).⁴²

Table 4 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 (reported in the fifth column). In the double removal, because Peanut M&M customers can no longer buy Snickers, and Snickers customers can no longer buy Peanut M&Ms, Mars bears 96.7% of the cost of the stockout. In the Snickers removal, most of the cost appears to be born by the downstream firm; one potential explanation is that among consumers who choose another product, many select another Mars Product (Twix or Peanut M&Ms). We also see the impact of each product removal on other manufacturers. Hershey (which owns Reese's Peanut Butter Cups

⁴⁰We report detailed product-level results from the joint removal in table 16 of the appendix (section A.5); product-level results from the two single-product removals are described in Conlon and Mortimer (2010). The product-level results suggest that the rebates may lead Mark Vend to stock products that compete less strongly with the major Mars brands.

⁴¹Reese's Peanut Butter Cups were not stocked by Mark Vend during either of the single-product removals, and so it does not appear as a top five substitute in those results.

⁴²Total losses appear smaller in the double-product removal in part because we sum over a smaller sample size of viable machine-treatment weeks (89) for this experiment, compared to the Peanut M&Ms removal (with 115 machine-treatment weeks).

and Hershey’s Chocolate Bars) enjoys relatively little substitution in the Snickers removal, in part because Reese’s Peanut Butter cups are not available as a substitute. In the double removal, when Peanut Butter Cups are available, Hershey profits rise by nearly \$61.43, capturing about half of Mars’ losses. We see substitution to the two Nestle products in the Snickers removal, so that Nestle gains \$19.32 as consumers substitute to Butterfinger and Raisinets; Nestle’s gains are a smaller percentage of Mars’ losses in the other two removals.

Direct analysis of the product removals can only account for the marginal cost aspect of the rebate (i.e., the price reduction given by Δ); one requires a model of restocking in order to account for the threshold aspect, $\overline{q_M}$. By more evenly allocating the costs of stocking out, the rebate should better align the incentives of the upstream and downstream firms, and lead the retailer to increase the overall service level. Similar to a two-part tariff, the rebate lowers the marginal cost to the retailer and reduces the margin of the manufacturer. Returning to table 3, the right-hand panel reports the retailer’s profit loss from the product removals after accounting for his rebate payments, assuming he qualifies. We see that the rebate reallocates approximately (\$17, \$30, \$50) of the cost of the Snickers, Peanut M&Ms, and joint product removals from the upstream to the downstream firm. The last column of table 4 shows that after accounting for the rebate contract, the manufacturer bears about 50% of the cost of the Peanut M&Ms removal, 60% of the cost of the joint removal, and 12% of the cost of the Snickers removal.

4 Estimation

4.1 Demand

The intuition provided in our theoretical framework is that the welfare effects of the vertical rebate depends on a few critical inputs. Those are: the substitutability of products in the downstream market, how the costs of reduced effort are distributed across the supply chain, and whether or not effort acts as a substitute or a complement in the profit function of upstream manufacturers. In order to consider the optimal product assortment, we need a parametric model of demand that predicts sales for a variety of different product assortments. We consider two such models: the nested logit and the random-coefficients logit, which are estimated from the full dataset (including both experimental and non-experimental periods). We discuss the random-coefficient model here; discussion and results from the nested-logit specification are contained in section A.6 of the appendix.

We consider a model of utility in which consumer i receives utility from choosing product

j in market t of:

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt}. \quad (2)$$

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.

A random-coefficients logit specification allows for correlation in tastes across observed product characteristics.⁴³ 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 consumers 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:

$$p_{jt}(\delta, \theta, a_t) = \int \frac{e^{\delta_{jt} + \sum_l \sigma_l \nu_{ilt} x_{jl}}}{1 + \sum_{k \in a_t} e^{\delta_{kt} + \sum_l \sigma_l \nu_{ilt} x_{kl}}} f(\nu_{ilt}|\theta). \quad (3)$$

We define a_t as the set of products stocked in market t , and a market as a machine-visit pair (i.e., a_t is the product assortment stocked in a machine between two service visits).⁴⁴ We specify $\delta_{jt} = d_j + \xi_t$; that is, we allow for 73 product intercepts as well as market-specific demand shifters. We allow for three random coefficients, corresponding to consumer tastes for salt, sugar, and nut content.⁴⁵ We estimate the parameters of the choice probabilities via maximum simulated likelihood (MSL). The log-likelihood is:

$$l_t(\mathbf{y}_t | \delta_{jt}, \theta, a_t) \propto \sum_j y_{jt} \log p_j(\delta_{jt}, \theta, a_t). \quad (4)$$

⁴³See Berry, Levinsohn, and Pakes (1995).

⁴⁴There are virtually no ‘natural’ stock-outs in the data; thus, changes to product assortment happen for two reasons: (1) Mark Vend changes the assortment when re-stocking, or (2) our field experiment exogenously removes one or two products.

⁴⁵Nut content is a continuous measure of the fraction of product weight that is attributed to nuts. We 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.

where y_{jt} are sales of product j in market t .⁴⁶

We report the parameter estimates in table 5. We report two levels of aggregation for ξ_t . The first allows for 15,256 fixed effects, at the level of a machine-service visit, while the second allows for 2,710 fixed effects, at the level of a machine-choice set. We report the log-likelihood, Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) for each specification. We use BIC to select the specification with 2,710 fixed effects.

4.2 Dynamic Model of Re-stocking

A crucial question for empirically evaluating the welfare effects of vertical rebate contracts is the size of the efficiency gains that are due to greater retailer effort, relative to the potential effects of foreclosure. Comparing the two effects is often challenging, because one needs to understand how effort endogenously responds to different contractual forms and product assortments. In most empirical contexts, the econometrician has very little data on the cost of effort. In this section we consider the specific case in which the retailer chooses the restocking frequency.

Our approach follows Rust (1987), but ‘in reverse.’ Rather than assuming that observed retailer wait times are optimal and using Rust’s model to estimate the cost of re-stocking, we use an outside estimate of the cost of re-stocking based on wage data from the vending operator, and compute the optimal wait time until the next restocking visit from the model. Specifically, we consider a multi-product (\mathbf{s}, \mathbf{S}) policy, in which the retailer 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 retailer’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 individual product when making restocking decisions. This closely parallels the problem of Mark Vend.⁴⁷

⁴⁶As in previous work, we do not estimate a price coefficient because there is no price variation in our data to identify the parameter. See Conlon and Mortimer (2013a) for a discussion of this issue.

⁴⁷That is, Mark Vend has information on whether particular days are likely to be busy or not, but does not observe the actual inventory levels of individual products until visiting the machine to restock it. In other retail contexts this assumption might be less realistic and could be relaxed; its role is primarily to reduce the computational burden in solving the re-stocking problem.

Mark Vend solves the following dynamic stocking problem, where $u(x)$ denotes the cumulative variable retailer profits after x potential consumers have arrived. Profits are not collected by Mark Vend until he restocks. His value function is:

$$V(x) = \max\{u(x) - FC + \beta V(0), \beta E_{x'}[V(x'|x)]\}. \quad (5)$$

The problem posed in (5) 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 pay-off \tilde{u} , defined as:

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

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

$$V(x, x^*) = (I - \beta \tilde{P}(x^*))^{-1} \tilde{u}(x, x^*). \quad (6)$$

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 $(u^R(x) + u^M(x))$, which incorporates the profits of the dominant upstream manufacturer. Likewise, we can consider the industry-optimal policy by replacing $u(x)$ with $(u^R(x) + u^M(x) + u^H(x) + u^N(x))$.

To find the optimal policy we iterate between (6) and the policy improvement step:

$$x^* = \min x : u(x) - FC + \beta V(0, x^*) \geq \beta P(x'|x)V(x', x^*). \quad (7)$$

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

In order to compute the dynamic restocking model, we use the following procedure. To obtain $u(x)$, we use the demand system generated by the random-coefficients model to simulate consumer arrivals and update inventories accordingly. We use actual machine capacities for each product.⁴⁸ We simulate 100,000 chains of consumer arrivals and construct

⁴⁸These capacities are nearly uniform across machines, and are: 15-18 units for each confection product,

the expected profit after X consumers arrive. We define our state variable to be the number of consumers expected to make a purchase from a hypothetical ‘full machine,’ and we define a ‘full machine’ as one that contains a set of 29 products, which we report in table 6, plus all of the products in the confections category.⁴⁹

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 incremental mileage. We convert 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.’ We then fit a smooth Chebyshev polynomial, and use this as our approximation of accumulated variable profits, $\hat{u}(x)$.⁵⁰

There is some heterogeneity in both the arrival rate of consumers to machines, as well as the service level of different machines in the data, so we divide our sample into four groups of machines 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 7. Our counterfactual analyses are based on cluster D, which is the largest cluster, containing 28 of the 66 machines in our sample. Machines in clusters A and C are smaller in size, while the seven machines 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. The distribution of daily sales for the machines in cluster D determine 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 assume a fixed cost of a restocking visit, $FC = \$10$, which approximates the per-machine restocking cost using the driver’s wage and average number of machines serviced per day. As a robustness test, we also consider $FC = \{5, 15\}$, which generate qualitatively similar predictions. In theory, one 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 a day before

11-12 units for each salty snack product, and around 15 units for each cookie/other product.

⁴⁹A typical machine in our dataset holds fewer products than this.

⁵⁰We designate our state space in terms of ‘expected sales at 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 X and the state space x is well defined, because $x \sim Bin(X, 1 - s_0)$ 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 $R^2 \geq 0.99$.

any products have stocked-out.⁵¹ This is helpful as an experimental control, but makes identifying FC from data impossible.⁵²

5 Results

5.1 Endogenous Effort

We now consider the results of the model in which we allow the re-stocking policy to endogenously respond to the wholesale prices (w_M, w_H, w_N) and the AUD contract (Δ, \bar{q}) . We begin by analyzing the retailer’s choice of effort, conditional on product assortment. For this analysis we construct a representative machine for which demand is described by the random-coefficients model from table 5, and the arrival rate of consumers is described by the process from the previous section, drawing from the distribution of daily sales for machines in cluster D. We assume that the representative machine is stocked with the set of 29 base products of our hypothetical ‘full machine,’ including five confections products, plus two additional products from the confections category, which we allow to vary. We fix Mark Vend’s five most commonly-stocked chocolate confections products: four Mars products (Snickers, Peanut M&Ms, Twix, and Plain M&Ms), and Nestle’s Raisinets. We also assume that confections prices are the uniform \$0.75 per unit we observe in the data, and that manufacturer marginal cost is zero.⁵³ We allow the retailer to choose from six different products for the final two slots in the confections category: two Mars products (Milky Way and Three Musketeers), two Hershey products (Reese’s Peanut Butter Cup and PayDay), and two Nestle products (Butterfinger and Crunch).⁵⁴

⁵¹In conversations with the retailer about his service schedule, he mentioned two points. First, he suspected that he was over-servicing, and reduced service levels after our field experiment. Second, he explained that high service levels are important to obtaining long-term (3-5 year) exclusive service contracts with locations. These specific locations almost certainly do not reflect a company-wide servicing policy. Specifically, these are high-end office buildings with high service expectations. Public locations, such as museums and hospitals, have much higher levels of demand and higher rates of stock-out events. These public locations affect company-wide servicing policies, but are not good candidates for running a successful field experiment.

⁵²We do not consider possible dynamic considerations, in which 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). In other work, we find very little evidence that the subsequent consumer arrival rate is affected by the history of stock-outs.

⁵³The assumption of zero manufacturer marginal costs implies that any efficiency gains we estimate represent an upper bound, because higher manufacturer costs would reduce the upstream firm’s revenues from restocking, leading to smaller efficiency gains from increased downstream effort.

⁵⁴We do not have sufficient information on other products to consider them in our counterfactual analysis. For example, Hershey’s with Almonds is popular nationally, but is rarely stocked in our data. As a robustness test, we also consider the counterfactual in which Mark Vend substitutes for the five base confection products, and we try a third Mars product, Skittles. Mark Vend is always worse off in these cases, and for space concerns

We compute the optimal re-stocking policy under four variants of the profit function and report those results in table 8. The optimal policy is stated as the answer to the question “Re-stock after how many expected sales?,” so a lower number implies more frequent restocking (and higher cost) to the retailer. Consistent with industry practice, we assume that all products are restocked when the downstream retailer visits a machine. In the first variant of the profit function, we consider the policy that maximizes retailer profit at the pre-rebate wholesale prices (w_M, w_H, w_N); in the second variant, we consider the policy that maximizes retailer profit at the post-rebate prices ($w_M - \Delta, w_H, w_N$). We label these ‘Retailer-Pre,’ and ‘Retailer-Post.’ In the third variant, we consider the joint profits of the retailer and Mars, which we label ‘Integrated.’ For that case, w_M and Δ are irrelevant because they are merely a transfer between integrated parties. The policy of the vertically-integrated firm is important, because it provides information on the threshold \bar{q}_M . If Mars were perfectly informed about retail demand, it could choose the level of \bar{q}_M in order to maximize the bilateral surplus. Finally, we report the policy that would be optimal for the confections industry as a whole (i.e., it maximizes the joint profit of R, M, H , and N). We label this ‘Industry.’ Table 8 reports the optimal restocking policies for five of the fifteen (i.e., ‘six choose two’) possible product combinations. The remaining combinations are dominated for the retailer.

In the absence of the rebate, the retailer sets an effort level that is 9-11% too low compared to a vertically-integrated (Retailer-Mars) firm. Our experiment indicated that the marginal cost aspect of the rebate, Δ , shifts approximately 40% of the stock-out cost onto the retailer.⁵⁵ However, this appears to have modest effects on the retailer’s stocking policy, which increases by around 2% (or 20-25% of the effort gap). This implies that \bar{q}_M , the threshold, plays a larger role than the marginal cost reduction in enhancing downstream effort. When the two additional products are both Nestle or Hershey products, the vertically-integrated firm sets the highest stocking level (replacing after 232 or 237 consumers), and the gap between the retailer and the vertically-integrated firm’s incentives are largest, at about 11%. When both additional products are owned by Mars (Three Musketeers and Milky Way) the difference in incentives is smallest, at 8.6%. In other words, when both of the marginal products are owned by the vertically-integrated Mars-Retailer firm, the firm is willing to tolerate the higher levels of stock-out events that result from lower retailer effort, because substitute products are owned by the firm.

we do not report those results.

⁵⁵For example, table 4 reports that the retailer’s share of the cost of a stock-out increases from 14% to 50% for the Peanut M&Ms removal, and from 3% to 40% for the joint Peanut M&Ms/Snickers removal.

The industry-optimal policy (i.e., the policy that maximizes the joint profits of Retailer-Mars-Hershey-Nestle) might involve more or less effort than the vertically-integrated (Retailer-Mars) policy, depending on whether downstream effort acts as a substitute or a complement across different upstream firms. In this case, we find that near the optimum level of effort, the vertically-integrated $M - R$ firm would prefer a higher effort level than the Industry optimum, though only 1-2% higher (i.e., 244 versus 247 for the Reeses' Peanut Butter Cup - Three Musketeers assortment). In other words, downstream effort is substitutable across upstream firms, so that increased retailer effort lowers competitors' sales, rather than allowing competitors to free-ride on effort.

Figure 3 reports the variable profits of each firm, ignoring the fixed cost of restocking, as a function of the restocking policy. We report these profits for the case in which the retailer stocks Reese's Peanut Butter Cups (Hershey) and Three Musketeers (Mars) in the final two slots, and we denote this assortment as (H, M) . We see that for both Mars and the Retailer, profits are monotonically decreasing as downstream effort falls, or the expected number of consumers between restocking visits rises. However, for Hershey and Nestle, profits initially increase as downstream effort falls. This happens because a few key Mars products sell-out faster than the Hershey and Nestle products, so that Hershey and Nestle benefit from forced substitution by consumers who arrive to the machine after the Mars products have sold out. Once effort falls below 400 expected sales, Hershey's and Nestle's profits increase with downstream effort, in line with Mars and the Retailer. At these low levels of service, downstream effort is complementary across upstream firms, so that all upstream firms benefit from greater downstream effort.

The optimal policies identified in table 8 are always less than 270, implying that over the relevant part of the curve, downstream effort is substitutable, and thus, increased retailer effort has a negative impact on Hershey and Nestle. As a result, if Mars chose \bar{q}_M to perfectly implement the vertically-integrated level of effort, it may lead to an over-provision of effort from the industry perspective (though not necessarily from a social perspective, because consumers benefit from the reduced stock-out events that result from greater retailer effort). The second implication of the results in figure 3 is that neither Hershey nor Nestle have an incentive to induce retailer effort through the use of a vertical rebate as Mars does. This is consistent with the fact that neither company does so in reality.⁵⁶

⁵⁶Hershey offers a 'rebate,' but it is equivalent to a reduced wholesale cost because the sales threshold and facing requirement are both zero.

5.2 Effects of the AUD on Product Assortment

Our next goal is to understand the relationship between the contractual structure and the retailer’s endogenous choice of product assortment. One obvious impact of the contract for product assortment occurs through the facing requirement. Four of our five baseline products are produced by Mars. A six-product facing requirement implies that both of the two additional products stocked by Mark Vend would need to be Mars products in order for him to qualify for the rebate payment. However, facing requirements can be difficult to enforce, and are rarely used in the absence of more easily enforced requirements such as a quantity target. Thus, it is instructive to examine the potential effect that the discount and quantity target have for Mark Vend’s assortment decision.

In order to do this, we consider three of the downstream re-stocking policies from table 8: Retailer-Pre, Retailer-Post, and Integrated. We do not consider the Industry-optimal policy, because there is no credible way to implement it with bilateral vertical arrangements. We compute profits throughout the supply chain for each of the three re-stocking policies using the five product assortments in table 8. In table 9, we report a subset of the three most relevant product assortment choices using the long-run expected profit from a single machine in cluster D (our ‘above-average’ group).⁵⁷

5.2.1 Role of the Discount

Table 9 shows that at the observed wholesale prices (w_M, w_H, w_N) ignoring any rebate payments, the retailer would choose to stock two Hershey products in the final two slots: Reese’s Peanut Butter Cup and Payday, which we refer to as the (H,H) assortment. This is illustrated by comparing across the three ‘Retailer-Pre’ rows to choose the assortment that maximizes profits in the ‘Retail No Rebate’ column.⁵⁸ This outcome is obtained because the retail margin is higher on Hershey products (i.e., retail prices are the same for both products, but $w_H < w_M$), and despite the fact that the Hershey products achieve slightly lower sales than the Mars products.⁵⁹

If we assume prices are fixed at the post-rebate levels for all product assortments ($w_M -$

⁵⁷Specifically, we report results for the three assortments for which at least one variant of the payoff function is maximized. For example, the choice of the two Nestle products (Butterfinger and Nestle Crunch) in the final two slots is never preferred by the retailer to the two better-selling Hershey products for any variant of the payoff function.

⁵⁸The Retailer-Pre row that corresponds to the retailer’s optimal assortment of (H,H) is highlighted.

⁵⁹This is confirmed by examining the columns that report total Industry profits and consumer surplus, which are higher for both the Three Musketeers - Milky Way, or (M,M) and (H,M) assortments.

Δ, w_H, w_N), with no quantity threshold or facing requirement, then the retailer would choose the Retailer-Post level of service, and he would earn $\$36,086 + \$2,096 = \$38,182$ for stocking the two Mars products (M,M) (Milky Way and Three Musketeers), $\$36,656 + \$1,617 = \$38,273$ for stocking (H,H), and $\$36,394 + \$1,882 = \$38,276$ for stocking an (H,M) assortment (Reese's Peanut Butter Cups and Three Musketeers). Thus, in the absence of any quantity threshold, the retailer would stock the (H,M) assortment. This would be a bad outcome for Mars, because it would earn only $\$11,763 - \$1,882 = \$9,881$ under this assortment and effort level, whereas under the Retailer-Pre assortment and effort in which the retailer chooses (H,H) and Mars does not pay a rebate, it earns $\$10,055$. Thus, the rebate is not individually rational for Mars in the absence of the quantity threshold or facing requirement, and should not be offered in equilibrium.⁶⁰

5.2.2 Role of the Facing Requirement and Quantity Threshold

If Mars sets a facing requirement of six products, then the retailer chooses a Retailer-Post level of service, and stocks both Mars products (M,M). This may be verified by comparing the Retailer-Post rows in table 9. Under this regime, the retailer earns $\$36,086 + \$2,096 = \$38,182$ by stocking (M,M) instead of the $\$36,661$ he earns by stocking (H,H) and not collecting the rebate (i.e., the rebate provides a $\$1,521$ increase in Retailer profit compared to the Retailer-Pre outcome).⁶¹ Thus, the rebate leads to foreclosure of both Hershey products.

In the absence of the facing requirement, Mars can alternatively use the quantity threshold to achieve the same retail assortment. To illustrate, suppose that Mars sets a sales threshold so that the rebate is only paid if Mars makes more than $\$11,912$ in revenue.⁶² This has the same effect as using a facing requirement because the retailer only obtains the rebate payment when he stocks (M,M).

Given the choice of an (M,M) assortment, Mars can also use the quantity threshold aspect of the rebate to push the retailer to a higher level of service. This may be seen in table 9 in the row associated with the Integrated stocking policy for the (M,M) assortment. At a threshold of $\$13,195$ in Mars revenues, the retailer would increase his service level to the

⁶⁰For this reason, the Retailer-Post row for the (H,M) assortment is not highlighted in table 9.

⁶¹One may work directly with Mars revenue in this example rather than \bar{q}_M because the retail and wholesale prices are uniform and we have assumed a zero marginal cost of production. One also needs to confirm that offering the rebate is individually rational for Mars. It is, as we describe in the next subsection.

⁶²A threshold of $\$11,912$ reflects Mars' profits under a level of effort for the (H,M) assortment that just exceeds the vertically-integrated effort. One needs this threshold instead of the Retailer-Post Mars revenues of $\$11,763$ under the (H,M) assortment to prevent the retailer from increasing his effort level slightly while still choosing the (H,M) assortment.

Integrated policy of 243 expected sales, and would earn $\$36,035 + \$2,111 = \$38,146$, which is better than the profit that he attains by forgoing the rebate, providing the Retailer-Pre level of service, and choosing the (H,H) assortment.

These effects on product assortment and effort are illustrated in figure 4, which plots the post-rebate retail profits under the two relevant product assortments, (H,M) and (M,M). The post-rebate retail profits are plotted against Mars' profits on the horizontal axis. The curve on the left, which is associated with lower levels of profit for Mars, reflects the retailers' post-rebate profits under the (H,M) assortment. Two red marks on the curve illustrate two of the stocking policies reported in table 9: the red mark at the apex of the profit function reflects the optimal service level at post-rebate prices; the red mark to the right of the apex reflects the service level that would be optimal from the perspective of a vertically-integrated Mars-Retailer firm. The product assortment associated with higher retail profit is the (H,M) assortment (i.e., the curve on the left attains higher retail profits than the curve on the right). However, Mars can use the threshold to induce the retailer to change from the (H,M) assortment to the (M,M) assortment. The vertical line that intersects the post-rebate retail profit curve (just to the right of the red mark indicating the optimal level of effort for the vertically-integrated M-R firm) is the point at which the retailer wants to switch from the (H,M) profit curve on the left to the (M,M) profit curve on the right. Thus, a threshold between \$11,912 and \$13,101 induces the retailer to drop the Hershey product in favor of the second Mars product. Once the retailer has switched to the (M,M) product assortment, Mars can push the retailers' service level (and Mars' profits) up a bit more by increasing the threshold to \$13,195, which is the level of retailer effort that is optimal for the vertically-integrated M-R firm.⁶³

In table 10, we summarize the potential competitive responses to the Retailer-Post and Integrated quantity thresholds under an (M,M) assortment. We compare both outcomes against the Retailer's optimal choice in the absence of the Mars rebate contract, which is an assortment choice of (H,H) at a Retailer-Pre level of effort, holding all other wholesale prices fixed.

The first alternative we examine is the (M,M)/Retailer-Post assortment, in which Mars' AUD payments would require sales of Mars products to exceed the threshold of 11,912 identified in figure 4. By stocking (M,M) and providing the 'Retailer-Post' level of effort, the Retailer's profits increase by \$1,521 ($\$38,182 - \$36,661$ from table 9). Mars' revenues increase

⁶³The retailer's profit curve associated with an (H,H) assortment would lie below and to the left of the (H,M) curve, and would not affect Mars' choice of the quantity threshold. Retail profits under the (H,H) assortment could potentially affect assortment if his binding constraint is the facing requirement.

from \$10,055 (the (H,H) assortment at Retailer-Pre effort level) to $\$13,101 - \$2,096 = \$11,005$, which is an extra \$950 for Mars. Thus, the rebate contract increases overall bilateral surplus by \$2,471 (\$1,521 for the retailer and \$950 for Mars). However, because Hershey's loss (of \$2,173) exceeds the retailer's gain, Hershey could avoid exclusion by lowering its wholesale price to $c_w = 12.83$ cents per unit from its current level of around 40 cents, thus transferring \$1,521 to the retailer. We do not have hard evidence on manufacturing costs in the industry, but we believe them to be in the neighborhood of 15-20 cents per bar. If so, Hershey would not be able to avoid exclusion once manufacturing costs were taken into account. Such foreclosure may be positive or negative from a social welfare perspective. Nestle loses \$31 under the new retail assortment and effort levels, so the rebate appears to increase overall profit for the industry by \$267.⁶⁴ Consumer utility is also higher by 291 units.

The second alternative sets a higher threshold (equal to 13,195) in order to induce the higher, vertically-integrated level of retailer effort. At this higher threshold, the Retailer's profit increases by \$1,485 relative to the non-rebating (H,H) option, and Mars' profits increase by \$1,029. Hershey still loses \$2,173, implying that they could avoid exclusion with a wholesale price of 13.54 cents per unit, holding fixed all of Mars' contractual terms. Overall industry profits again exceed those under the (H,H)/Retailer-Pre outcome, by \$302 for the industry, and consumer utility increases by 444 units.

These results suggest that Hershey would not be able to avoid exclusion by outbidding Mars for placement in the retail assortment. One may conversely ask whether Mars could reduce its discount and still obtain exclusivity. The last row in table 10 provides insight into this question. In this row, we assume that each manufacturer's production cost is \$0.15 per unit, so that Hershey's best offer to the retailer is a wholesale price of \$0.15. Assuming that Hershey lowers its wholesale price to \$0.15, we compute the per-unit transfer that Mars must make to the Retailer to make him indifferent between the (H,H) assortment at Hershey's best price of \$0.15 and the (M,M) assortment under Mars' current facing and quantity threshold requirements. We report this transfer as a percentage change from Mars' current discount, Δ . Under the lower Retailer-Post threshold, Mars could reduce the generosity of its discount by 5.3% and maintain exclusivity. Under the higher Integrated threshold, the discount may be reduced by 3.5%. Assuming that the \$0.15 production cost estimate is reasonable, this gives some indication that the terms of Mars' current rebate program are well designed.

⁶⁴The change in industry profit is the \$2,471 Retailer-Mars gain, less Hershey's \$2,173 and Nestle's \$31 losses.

5.3 Pure Efficiency Effect

Setting aside the potential competitive effects of the contract, one may also examine its pure efficiency effect. We define the pure efficiency effect as the mitigation of downstream moral hazard and inducement of additional downstream effort. To quantify this effect, we hold assortment fixed, and measure the welfare impact of moving from the ‘Retailer-Pre’ row to the ‘Integrated’ row in table 9.⁶⁵ This is likely to represent an upper bound on the efficiency effect because it assumes a zero marginal cost of production upstream.⁶⁶

We report these efficiency calculations in table 11 for the three product assortments (M,M), (H,M), and (H,H). The first row reports the policy change that would result from moving from the Retailer-Pre effort level to the Integrated level of effort (i.e., the same ‘Retail Understock’ from table 8). For the case of (M,M), the likely result under the AUD, Retailer’s profits decline (from increased effort) by \$55, Mars’ revenues rise by \$128, and Nestle’s profits fall by \$10 (due to less forced substitution from stocked-out Mars products to Raisinets. This results in an industry-wide change to producer surplus of \$64. Consumer utility rises by 209 units, which represents a gain of roughly one percent relative to the consumer surplus that is achieved without the additional effort induced by the AUD. Table 11 shows similar efficiency gains using the same calculations for the two other product assortment choices ((H,M) and (H,H)) from table 9. Thus, conditional on the product assortment, potential efficiencies from increased effort are small, and primarily accrue to consumers rather than firms. This is due to the fact that retail prices do not adjust, and is in spite of the fact that the potential impact of the contract on Retailer effort is substantial, requiring increases in the re-stocking rate of 8.6-11% in order to achieve the vertically-integrated outcome.

5.4 Comparison to Linear Pricing

The results reported so far hold fixed the wholesale prices of all three upstream firms. While allowing for endogenous wholesale pricing is outside the scope of our model, we can construct an important alternative baseline comparison, which is to allow Mars to re-optimize its wholesale price in the absence of an AUD.⁶⁷ We consider a subgame perfect solution for which, in the first stage, Mars sets w_M , and in the second stage, the retailer chooses a product assortment a and an effort level or restocking frequency e . We hold (w_H, w_N) fixed

⁶⁵The vertically-integrated level of downstream effort is a convenient baseline for comparison. Recall that as long as effort acts as a substitute upstream, Mars’ profits rise more quickly than the cost of the rebate.

⁶⁶We use the \$10 cost per restocking visit for the retailer.

⁶⁷One can interpret this as an AUD contract where the threshold $\bar{q} = 0$.

and examine Mars' unilateral response.⁶⁸ We consider this exercise more like a comparative static rather than a counterfactual equilibrium.

Table 12 provides details on the outcome of this linear pricing approach. The first two columns provide the benchmark restocking policies and profit/surplus numbers from table 9, and the last column reports the same figures for Mars' optimal linear wholesale price. Holding competitors' prices fixed, but allowing for endogenous retailer decisions for assortment and effort, Mars undercuts Hershey's prices in order to obtain the exclusive (M,M) assortment.⁶⁹ The resulting linear price of 41.36 cents per unit is more than 15% lower than the observed post-rebate price. This lower price induces greater retail effort than the post-Rebate prices (ignoring the facing and quantity threshold requirements), resulting in a retailer restocking policy of waiting for 257 expected sales, instead of 259. However, the lower linear price does not reach the vertically-integrated optimal rate of restocking, which is after 243 expected sales.⁷⁰ We hold retail prices fixed, so consumer utility is affected only by changes in the assortment and retailer effort level, and not the wholesale prices directly. Consumers are slightly better off with the effort and assortment decisions that arise under linear wholesale pricing than they are under the post-rebate wholesale pricing, but not as well off as they are under the vertically-integrated rebate outcome (which can be attained with an Integrated choice of \bar{q}_M under the AUD). From a social surplus perspective, the costs of banning the AUD and restricting Mars to linear contracts appear to be small, especially if it is difficult to choose an appropriate threshold \bar{q}_M . For example, if Mars sets the sales threshold too low, the linear contract outperforms the AUD contract both for the industry as a whole, and for consumers. However, consistent with some of the existing literature (see Kolay, Shaffer, and Ordover (2004)) the AUD has the potential to shift substantial rents from retailers to upstream firms (in this case, almost 10% of Mars' revenue).

⁶⁸A potential justification might be if retailers could purchase Hershey or Nestle products through some other channel at prevailing wholesale prices, such as through a warehouse club like Costco. One might also want to allow Hershey or Nestle to adjust their prices in response to Mars' price. We again refer the reader to section A.3 of the appendix, in which we show that this problem is very difficult, and need not lead to an equilibrium outcome because the discreteness of the retailer's product assortment decision leads to discontinuities in the reaction functions of the wholesalers.

⁶⁹Mars must provide more compensation to the retailer than $w_H - \epsilon$ because Reese's Peanut Butter Cups is preferred to Milky Way.

⁷⁰We also considered a deviation in which we reduced Hershey's price by 50% and solved for the endogenous Mars price, though for space considerations we do not report those results. In that case, and assuming that Mars is restricted to a single linear wholesale price on all products, it is no longer worthwhile for Mars to undercut Hershey. Instead, Mars responds by setting a higher linear wholesale price (i.e., a price between the pre- and post-rebate prices), at which price, the retailer selects (H,M) at a slightly lower level of effort than the post-rebate (H,M) level of 261.

5.5 Implications for Mergers

Vending is one of many industries for which retail prices are often fixed across similar products and under different vertical arrangements. Indeed, there are many industries for which the primary strategic variable is not retail price, but rather a slotting fee or other transfer payment between vertically-separated firms. Thus, our ability to evaluate the impact of a potential upstream merger may turn on how the merger affects payments between firms in the vertical channel.

In this analysis, we consider the impact of three potential mergers (Mars-Hershey, Mars-Nestle, and Hershey-Nestle) on the AUD terms offered to the retailer by Mars. Given the degree of concentration in the confections industry, antitrust authorities would likely investigate proposed mergers, especially mergers involving Mars.⁷¹ In order to analyze the impact of any potential merger, we conduct a similar exercise as before, but consider the incentives of the merged firm.

Table 13 considers the assortment and effort decisions of the retailer after a hypothetical Mars-Hershey merger. We assume that the Hershey product (Reeses Peanut Butter Cup) is priced at the Mars wholesale price and included in Mars' rebate contract after the merger. The merged (Mars-Hershey) firm is now happy for consumers to substitute to Reese's Peanut Butter Cups, and the AUD is able to achieve the industry-optimal (and socially-optimal) product assortment of (H,M). The merged firm faces competition from Nestle (Crunch and Butterfinger), which charges lower wholesale prices.⁷² In the absence of an AUD, the Retailer maximizes profits by stocking the two Nestle products (earning \$36,594), but the AUD induces the retailer to choose (H,M), as well as the effort level that would be set by the vertically-integrated firm (earning $\$36,340 + \$2,105 = \$38,445$). Section A.7 of the Appendix provides details and discussion on two additional mergers: Mars-Nestle and Hershey-Nestle. A pattern similar to the results from table 9 emerges in both cases. Specifically, the retailer's optimal assortment in the absence of an AUD is (H,H), but the AUD induces the retailer to stock (M,M). The (M,M) assortment attains greater industry and social surplus, but does not achieve the industry- or socially-optimal assortment of (H,M).

We summarize the impact of all three mergers for the AUD contract in Table 14. We focus on the extent to which Mars could reduce the generosity of the rebate, holding retail prices fixed and assuming a 15 cent marginal cost of production for all upstream firms. The first column reports outcomes in the absence of any mergers, replicating the second

⁷¹For a related analysis of diversion ratios in this market, see Conlon and Mortimer (2013b).

⁷²We assume Nestle's wholesale prices do not adjust after the Mars-Hershey merger.

column of table 10 The second column reports outcomes after a hypothetical merger of Mars and Hershey. In this case, the merged Mars-Hershey firm is able to implement the industry-optimal product assortment of (H,M). Holding fixed the terms of the rebate, this merger appears to be an unambiguous improvement for consumers, the industry, and to the merging parties. However, since Nestle is less strong as a competitor, and because (H,M) is more profitable than (M,M), Mars may reduce the generosity of the rebate significantly, lowering the discount by as much as 42.3%, assuming a 15 cent marginal cost for the industry.

The third column reports outcomes after a hypothetical Mars-Nestle merger. The primary effect of this merger is that Mars is able to control additional profits from Raisinets. This means that Mars can transfer the same amount of profits to the retailer with a smaller rebate (in percentage terms). Assuming a 15 cent marginal cost, the merged firm could reduce the rebate by roughly 12%. Finally, the last column provides outcomes under a hypothetical Hershey-Nestle merger. The primary impact in this case is that the merged firm has control of more surplus when trying to avoid exclusion. However, the profit loss is relatively unaffected, so it appears unlikely to substantially impact the rebate paid to the retailer.

Throughout the paper, we report the variable profits for the retailer; it is likely that his overall operating profits after accounting for administrative and overhead costs, are substantially lower. In the *Intel* case, the rebate program was reported to account for more than one quarter of Dell's operating profits. Based on communication with industry participants, we think that the Mars rebate may be an even larger fraction of operating profits in the vending industry. This means that a 42% rebate reduction (implied by the hypothetical Mars-Hershey merger) may represent a substantial fraction of the overall operating profits of the retailer.

6 Conclusion

Using a new proprietary dataset that includes exogenous variation in product availability, we provide empirical evidence regarding the potential efficiency and foreclosure aspects of an AUD contract. Similar vertical rebate arrangements have been at the center of several recent large antitrust settlements, and have attracted the attention of competition authorities in many jurisdictions.

In order to understand the relative size of the potential efficiency and foreclosure effects of the contract, our framework incorporates endogenous retailer effort and product assortment decisions. A discrete-choice demand model allows us to characterize the downstream

substitutability of competing products, and combining this with a model of retailer effort allows us to estimate the impact of downstream effort across upstream and downstream firms. Identification of both the demand and retailer-effort models benefit from exogenous variation in product availability made possible through a field experiment. We show that the vertical rebate we observe has the potential to increase effort provision by roughly 9-11%, but these rents are mostly captured by consumers. The rebate also enables the dominant firm, Mars, to foreclose Hershey by leveraging profits from dominant brands such as Snickers and Peanut M&Ms, and to obtain shelf-space for brands such as Milky Way.

We find that at the prevailing wholesale prices, this foreclosure enhances the profitability of the overall industry and improves social surplus, but does not lead to a product assortment that maximizes industry profits. We note that in the absence of the vertical rebate, manufacturers may charge different wholesale prices. In a limited comparison of Mars' optimal linear wholesale prices to the AUD contract, we find that the primary difference between Mars' AUD and linear wholesale pricing is the allocation of profits between the dominant upstream firm and the retailer. The differential impact on social welfare is small, and depends on how the dominant firm sets the quantity threshold in the AUD. Finally, we explore the potential impact of three potential upstream mergers on the likely terms of the AUD contract, holding retail prices fixed. We find that a merger between the two largest upstream firms has the potential to induce the socially-optimal product assortment, but may also lead to a reduction in the rebate payments to retailers.

In addition to providing a road-map for empirical analyses of vertical rebates, and results on one specific vertical rebate, our detailed data and exogenous variation allow us to contribute to the broader literature on the role of vertical arrangements for mitigating downstream moral hazard and inducing downstream effort provision. Empirical analyses of downstream moral hazard are often limited not only by data availability, but also by the ability to measure effort, and our setting proves a relatively clean laboratory for measuring the effects of downstream effort.

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).
- BAKER, G. P., AND T. N. HUBBARD (2003): “Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information,” *American Economic Review*, 93(3), 551–572.
- BERNHEIM, B. D., AND M. WHINSTON (1998): “Exclusive Dealing,” *Journal of Political Economy*, 106(1), 64–103.
- BERRY, S. (1994): “Estimating discrete-choice models of product differentiation,” *RAND Journal of Economics*, 25(2), 242–261.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- 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.
- (2014): “All-Units Discounts as a Partial Foreclosure Device,” Working Paper.
- CHEN, C.-W. (2014): “Estimating the Foreclosure Effect of Exclusive Dealing: Evidence from the Entry of Specialty Beer Producers,” Working Paper.
- CHEN, Z., AND G. SHAFFER (2014): “Naked Exclusion with Minimum-Share Requirements,” *Rand Journal of Economics*, 45(1), 64–91.
- CHIPTY, T. (2001): “Vertical Inegration, Market Foreclosure, and Consumer Welfare in the Cable Television Industry,” *American Economic Review*, 91(3), 428–53.
- 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.
- CRAWFORD, G., R. LEE, M. WHINSTON, AND A. YURUKOGLU (2015): “The Welfare Effects of Vertical Integration in Multichannel Television Markets,” Working Paper.
- 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.
- ELHAUGE, E., AND A. WICKELGREN (2012): “Anticompetitive Market Division through Loyalty Discounts without Buyer Commitment,” Working Paper.
- (2014): “Robust Exclusion and Market Division Through Loyalty Discounts with Buyer Commitment,” Working Paper.
- FIGUEROA, N., E. IDE, AND J.-P. MONTERO (2014): “Discounts as a Barrier to Entry,” Working Paper.
- FUMAGALLI, C., AND M. MOTTA (2006): “Exclusive Dealing and Entry, when Buyers Compete,” *American Economic Review*, 96(3), 785–795.
- GIL, R. (2007): “Make-or-buy in Movies: Integration and Ex-post Renegotiation,” *International Journal of Industrial Organization*, 25(4), 643–655.
- HO, J., K. HO, AND J. H. MORTIMER (2012a): “Analyzing the Welfare Impacts of Full-line Forcing Contracts,” *Journal of Industrial Economics*, 60(3), 468–498.
- (2012b): “The Use of Full-line Forcing Contracts in the Video Rental Industry,” *American Economic Review*, 102(2), 686–719.
- 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.
- HUBBARD, T. N. (1998): “An empirical examination of moral hazard in the vehicle inspection market,” *Rand Journal of Economics*, 29(2), 406–426.
- KLEIN, B., AND A. LERNER (2014): “The Economics of Alternative Antitrust Standards for Loyalty Contracts,” Working Paper.
- 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.
- LAFONTAINE, F., AND M. SLADE (2007): “Vertical Integration and Firm Boundaries: The Evidence,” *Journal of Economic Literature*, 45(3), 629–685.
- (2008): “Exclusive Contracts and Vertical Restraints: Empirical Evidence and Public Policy,” in *Handbook of Antitrust Economics*, ed. by P. Buccirossi, pp. 391–414. MIT Press, Cambridge, MA.
- LEE, R. S. (2013): “Vertical Integration and Exclusivity in Platform and Two-Sided Markets,” *American Economic Review*, 103(7), 2960–3000.
- 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.
- SASS, T. (2005): “The competitive effects of exclusive dealing: Evidence from the U.S. beer industry,” *International Journal of Industrial Organization*, 23(3-4), 203–225.
- 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.
- SINKINSON, M. (2014): “Pricing and Entry Incentives with Exclusive Contracts: Evidence from Smartphones,” Working Paper.
- STOKEY, N. L., R. LUCAS, AND E. PRESCOTT (1989): *Recursive Methods for Economic Dynamics*. Harvard University Press.


SUDHIR, K., AND V. R. RAO (2006): “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



Notes: The top panel illustrates a conventional non-linear pricing schedule, in which a retailer pays a wholesale price of 1 and receives a discounted price of 0.75 for quantities greater than 6. The bottom panel illustrates an All-Units Discount with a wholesale price of 1, a threshold equal to 6, and a discounted price equal to 0.67.

Figure 2: Mars Vend Operator Rebate Program

The Only Candy You Need To Stock In Your Machine!

Spinal#1	Spinal#2	Spinal#3	Spinal#4	Spinal#5	Spinal#6	Spinal#7	Spinal#8
#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!

6 **PrVend** 52 Weeks Ending 10/4/09 **MARS** chocolate north america

2010 Vend Operator Program

Platinum Rebate Level


- Receive a great Every Day Low Cost from your Authorized Vend Product Distributor
- Purchase brand level targets for 6 singles or king size items
 - Reduction from 7 must-stock items in 2009!
 - You pick the six items!
 - Will consolidate item variants to qualify (by brand, excluding SNICKERS® Bar and M&M's® Peanut Candies)
- No Growth Requirement
- PLUS a Rebate Payment **Low Cost PLUS Rebate:**

Item	Rebate %	Rebate \$ Per Bar (singles)
All Items	8%	4.0¢

11 **MARS** chocolate north america


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

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



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

Figure 4: Impact of AUD Quantity Threshold on Retail Assortment Choice



Notes: Figure reports retailer profit under two assortment choices ((H,M) on the left and (M,M) on the right), against sales of Mars products. For a quantity requirement of 11,912 or greater (noted by the vertical dashed line), the retailer prefers to switch his assortment from (H,M) to (M,M).

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

Manu- facturer	Product	National:			Mark Vend:	
		Rank	Avail- ability	Share	Avail- ability	Share
Mars	Snickers	1	89	12.0	96	22.0
Mars	Peanut M&Ms	2	88	10.7	96	23.0
Mars	Twix Bar	3	67	7.7	79	13.0
Hershey	Reeses Peanut Butter Cups	4	72	5.5	29	3.7
Mars	Three Musketeers	5	57	4.3	34	4.3
Mars	Plain M&Ms	6	65	4.2	47	6.4
Mars	Starburst	7	38	3.9	16	1.0
Mars	Skittles	8	43	3.9	77	6.5
Nestle	Butterfinger	9	52	3.2	33	2.7
Hershey	Hershey with Almond	10	39	3.0	0	0
Nestle	Raisinets	>45	N/R	N/R	78	8.9

Notes: National Rank, Availability and Share refers to total US sales for the 12 weeks ending May 14, 2000, reported by Management Science Associates, Inc., at <http://www.allaboutvending.com/studies/study2.htm>, accessed on June 18, 2014. National figures not reported for Raisinets because they are outside of the 45 top-ranked products. By manufacturer, the national shares of the top 45 products (from the same source) are: Mars 52.0%, and Hershey 20.5%. For Mark Vend, shares are: Mars 80.0%, Hershey 8.5% (calculations by authors). Mark Vend averages 6.86 confection facings per machine.

Table 2: Top 5 Substitutes (Based on Vends)

Snickers	Peanut M&Ms	Both
Peanut M&Ms*	Snickers*	Assorted Chocolate*
Twix Caramel*	Assorted Pretzel/Popcorn*	Twix Caramel*
Assorted Pretzel/Popcorn	Assorted Nuts*	Reese's Peanut Butter Cups*
Assorted Nuts*	Twix Caramel*	Assorted Pastry*
Assorted Fruit Snack*	Plain M&Ms*	Plain M&Ms*
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)

Notes: Asterisks indicate a statistically significant increase in sales during the treatment period (vs. the control period) at the 95% level.

Table 3: Downstream Profit Impact

			Before Rebate			After Rebate		
Exogenous Removal	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

Notes: Calculations by authors, using exogenous product removals from the field experiment.

Table 4: Upstream (Manufacturer) Profits

Exogenous Removal	% Born by Mars					
	Mars	Hershey	Nestle	Other	Before Rebate	After Rebate
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%

Notes: Calculations by authors, using exogenous product removals from the field experiment. The ‘% Born by Mars Before Rebate’ reports the percentage of the total cost of a product removal that is born by Mars, shout accounting for the rebate payment to the retailer. ‘% Born by Mars After Rebate’ is equivalently defined.

Table 5: Random Coefficients Demand Model

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

Notes: The random coefficients estimates correspond to the choice probabilities described in section 4, equation 3. Both specifications include 73 product fixed effects. Total sales are 2,960,315.

Table 6: Products Used in Counterfactual Analyses

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

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

Table 7: Summary of Sales and Revenues for Four Clusters of Machines

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

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

Table 8: Stocking Policies (Restock after how many expected sales?)

Additional Products Stocked*:	Retail	Retail	Integrated	Industry	Retail	Rebate
	-Pre	-Post			Understock	Increase
Reeses PB Cup (H), 3 Musketeers (M)	267	261	244	247	9.4%	2.2%
Reeses PB Cup (H), Payday (H)	263	257	237	241	11.0%	2.3%
3 Musketeers (M), Milkyway (M)	264	259	243	244	8.6%	1.9%
Reeses PB Cup (H), Butterfinger (N)	263	257	238	241	10.5%	2.3%
Butterfinger (N), Crunch (N)	257	251	232	234	10.8%	2.3%

Notes: The five base products are: Snickers, Peanut M&Ms, Twix, Plain M&Ms (Mars), and Raisinets (Nestle). Manufacturers are denoted as (H) for Hershey's, (N) for Nestle, and (M) for Mars. 'Retail understock' reports (Retail-Pre - Integrated)/Integrated.

Table 9: Profits under Alternate Product Assortments and Stocking Policies

Policy	Retail (No Reb.)	Rebate	Mars	Hershey	Nestle	Inte- grated	Industry	Consumer Utility
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers								
Retailer-Pre (267)	36,399	1,875	11,719	1,302	1,260	48,117	50,679	21,685
Retailer-Post (261)	36,394	1,882	11,763	1,299	1,257	48,157	50,713	21,752
Integrated (244)	36,335	1,899	11,871	1,290	1,249	48,206	50,744	21,911
(H,H) Assortment: Reeses Peanut Butter Cup and Payday								
Retailer-Pre (263)	36,661	1,609	10,055	2,173	1,285	46,716	50,174	20,845
Retailer-Post (257)	36,656	1,617	10,106	2,167	1,282	46,762	50,211	20,914
Integrated (237)	36,578	1,640	10,251	2,149	1,272	46,829	50,250	21,102
(M,M) Assortment: Three Musketeers and Milkyway								
Retailer-Pre (264)	36,090	2,091	13,067	0	1,256	49,156	50,412	21,080
Retailer-Post (259)	36,086	2,096	13,101	0	1,254	49,187	50,441	21,136
Integrated (243)	36,035	2,111	13,195	0	1,246	49,230	50,476	21,289

Notes: Profit numbers represent the long-run expected profit from a single machine in Group D. Each panel represents one of the product assortments in table 8.

Table 10: Potential Competitive Responses (Comparison is (H,H) Pre-Rebate)

Threshold	> 11,912 < 13,101	= 13,195
Assortment	(M,M)	(M,M)
Policy	Retailer-Post	Integrated
Δ Retail Profit	1,521	1,485
Δ Mars Revenue	950	1,029
Δ Hershey Revenue	-2,173	-2,173
Δ Nestle Revenue	-31	-39
Δ Retailer+Hershey	-652	-688
Δ Consumer Utility	291	444
Δ Industry Profit	267	302
Price to Avoid Exclusion	12.83	13.54
Rebate Reduction ($c = .15$)	5.27%	3.53%

Notes: Comparisons are to the retailer's 'Pre-Rebate' optimal stocking policy and (H,H) assortment. Results refer to the long-run impacts for a single machine in Group D. The small changes to Nestle's profit result from changes in the retailer effort level (i.e., Retailer-Post or Integrated effort vs. Pre-Rebate for the (H,H) assortment). Hershey's price is 42.75 cents per unit.

Table 11: Potential Efficiency Gains from Vertically-Integrated Stocking Decisions

	(M,M)	(H,M)	(H,H)
Policy Change	8.6%	9.4%	11.0%
Δ Retail	-55	-64	-83
Δ Mars	128	196	196
Δ Mars + Retailer	73	132	113
Δ Hershey		-12	-24
Δ Nestle	-10	-11	-13
Δ Producer Surplus	64	65	76
% Δ Consumer Utility	0.99%	1.04%	1.23%

Notes: Calculations based on holding assortment fixed and measuring the welfare impact of moving from the 'Retailer-Pre' row to the 'Integrated' row in table 9. Policy Change refers to the percentage change in restocking frequency that results from changing from a Retailer-Pre policy to an Integrated policy, and is also reported in table 8.

Table 12: Linear Pricing vs. AUD (Assortment is (M,M))

	Post-Rebate	Integrated	Linear Pricing
Threshold	>11,912	=13,195	=0
Restocking Policy	259	243	257
Retail Profit (including rebate)	38,182	38,146	39,103
Mars Profit	11,005	11,084	10,094
Nestle Profit	1,254	1,246	1,253
Consumer Utility	21,136	21,289	21,158

Notes: The optimal wholesale price under linear pricing is estimated to be 41.36 cents per unit. Hershey is excluded in the (M,M) assortment for all three arrangements, and earns zero profit. The small changes in Nestle's profits are due to the effect of changes in the retailer's choice of restocking policy on the sales of Raisinets.

Table 13: Profits after Mars-Hershey Merger

Policy	Retail (No Reb.)	Rebate	Mars/ Hershey	Nestle	Inte- grated	Industry	Consumer Utility
(H,M) Assortment: Reeses Peanut Butter Cup and Three Musketeers							
Retailer-Pre (267)	36,399	2,083	13,021	1,260	49,419	50,679	21,685
Retailer-Post (262)	36,395	2,089	13,055	1,257	49,451	50,708	21,741
Integrated (245)	36,340	2,105	13,155	1,249	49,496	50,745	21,903
(N,N) Assortment: Butterfinger and Crunch							
Retailer-Pre (257)	36,594	1,631	10,193	2,707	46,787	49,494	19,430
Retailer-Post (251)	36,589	1,639	10,246	2,700	46,835	49,535	19,502
Integrated (232)	36,514	1,662	10,386	2,681	46,900	49,581	19,687

Notes: Profit numbers represent the long-run expected profit from a single machine in Group D. The two panels represent the two product assortments that offer the greatest potential retailer profit under different stocking policies and rebate payments.

Table 14: Comparison under Alternate Ownership Structures

	No Merger	M-H Merger	M-N Merger	H-N Merger
AUD Assortment	(M,M)	(H,M)	(M,M)	(M,M)
Alternative	(H,H)	(N,N)	(H,H)	(H,H)
Policy	Integrated	Integrated	Integrated	Integrated
Δ Retail	1,485	1,851	1,689	1,485
Δ Mars	1,029	857	785	1,029
Δ Bilateral	2,514	2,708	2,474	2,514
Δ Competitor	-2,173	-1,458	-2,173	-2,212
Δ Retailer + Competitor	-688	393	-484	-727
Δ Consumer Surplus	444	2,473	436	444
Δ Industry	302	1,251	302	302
Price to Avoid Exclusion	13.54	n/a	9.52	14.05
Rebate Reduction ($c = 0.15$)	3.5%	42.3%	12.1%	2.3%

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

Appendix

A.1: Retailer Choice of Product 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 $e(a)$, so that 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 because 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 reduces to the “Knapsack Problem” and is NP-hard. Once we allow for substitutes or complements in the profit function, only brute-force solutions exist. This means that for each possible a , we must solve a dynamic programming problem for retailer effort and calculate profits. For the case in which the retailer chooses seven products to stock from a set of 12, or $\binom{12}{7}$, there are 792 possible combinations that must be considered.⁷³

In order to illustrate the role that an AUD plays in a retailer’s choice of product assortment, this section instead 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 effort decision 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. If the efficiency gains from an AUD contract are relatively small, this intuition should mimic the results of the dynamic model of retailer effort.

The intuition is captured by considering the situation in which M sells two products (products 1 and 3) and offers an AUD, and N sells a single product (product 2) under a simple linear price. We assume both upstream firms face production costs of zero, and sell

⁷³To 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 (Peanut M&Ms, Snickers, Twix, Plain M&Ms, 3 Musketeers, Milkyway, Skittles). We do not present results that consider removing Peanut M&Ms, Snickers, Twix, Plain M&Ms, or Raisinets from the choice set. In practice, we consider robustness to replacing Plain M&Ms or Raisinets, although 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. We report only the best alternatives among those 21 in table 8.

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.

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.⁷⁴

The choice variables are: w_M , and the threshold and discount terms for the AUD, \bar{q}_M and Δ , (for firm M), w_N (for firm N), 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_N 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 random ordering of consumers and the stochasticity of demand.

The results of the model are shown in Table 15. 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.

⁷⁴Random selection of consumers without replacement implies deterministic demand, although realized sales may vary based on arrival order and capacity.

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, while in the high-capacity case the exclusive arrangement maximizes industry profits about two-thirds of the time. Thus, whether or not the exclusion that results from the AUD maximizes industry profits depends on whether the impact of retail capacity on realized sales is complementary or substitutable across upstream firms.

A.2: Comparison with Other Contracts

One may consider 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 (LP), 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 expositional, 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 variable 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 $t(a, e)$:

$$\max_{a,e} \pi_R(a, e) - c(e) + p(a, e)t(a, e). \tag{8}$$

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:

$$p'(e)t(e) + p(e)t'(e) = \pi'_M(e).$$

One may 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 $t < 0$.⁷⁵ 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}_M 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, LP.⁷⁶ 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 $\pi'_M(e) - t'(e) > 0$.⁷⁷ The QF contract allows M to offer a contract that requires the integrated level of effort, through \bar{q}_M . 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}_M , which means the AUD is more flexible when there is uncertainty about downstream demand. In the absence of this uncertainty, the AUD mimics a QF contract.⁷⁸

⁷⁵The challenge of the 2PT is that the upstream firm M must determine the appropriate fixed fee $t(0)$. Kolay, Shaffer, and Ordover (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, in the absence of uncertainty an individually-tailored 2PT enables full extraction by M , but is a likely violation of the Robinson-Patman Act.

⁷⁶For the setting in which rebate contracts are not allowed and firms are required to offer linear wholesale prices, solving for optimal prices is difficult, because the solution depends both on the effort of the downstream retailer, and the endogenous product assortment, neither of which needs to be a smooth continuous function of wholesale prices. For this reason our empirical work considers deviations from observed prices rather than fully solving for a new equilibria in linear wholesale prices. The following section, A.3, provides further discussion.

⁷⁷This 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).

⁷⁸Chao and Tan (2013) explore connections between QF, AUD, and 3PT when a dominant manufacturer faces a capacity-constrained rival.

A.3: Solving for Endogenous Linear Prices

Allowing for endogenous linear prices (even in the absence of rebates) is a challenging extension. In our setting, the retail price is fixed, and the retailer chooses a product assortment by solving a binary integer programming problem. Even with two wholesale prices that are the same across all of a manufacturer’s products (w_M, w_N) , this is a challenging problem. The principle challenge is that small changes in one of the wholesale prices (e.g., w_M) can result in discrete changes in the product assortment, and discontinuous jumps in 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 (or none at all) may exist.

When we report Mars’ wholesale price under a linear contract in section 5, we hold the prices of competitors (Hershey and Nestle) fixed and solve for the optimal Mars price w_M . Even this simplified setting is challenging because Mars’ residual profit as a function of price takes on a sawtooth form. Mars’ profit increases with the wholesale price, until the retailer chooses to drop a Mars product in favor of a competitor’s product. There are potentially as many peaks on the sawtooth as there are dropped Mars products. This is demonstrated in figure 5.

A.4: Computing Treatment Effects

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

There are two challenges in implementing the removals and interpreting the data generated by them. The first challenge is that there is a large amount of variation in overall sales at the weekly level, independent of our exogenous removals. 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 product removals were

done during the summer of 2007, which was a high-point in demand at these sites, most likely due to macroeconomic conditions. In this case, using a simple measure like previous weeks' sales, or overall average sales for $E[q_{jt}^0]$ could result in unreasonable treatment effects, such as sales increasing due to product removals, or sales decreasing by more than the sales of the focal products.

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

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

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 the product removal, which would be problematic for our study.

We propose a slight modification of (9) which removes the bias. That is, we 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 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 the case that for some treatment weeks t the set of possible control weeks $S_t = \{\emptyset\}$. Under this definition of the control, some treatment weeks constitute ‘outliers’ and are excluded from the analysis. Of the 1470 machine-experiment-week combinations, 991 of them have at least one corresponding control week, and at the machine-experiment level, 528 out of 634 have at least one corresponding control. Each included treatment week has an average of 24 corresponding control weeks, though this can vary considerably from treatment week to treatment week.⁷⁹

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.

A.5 Product-level Results of Exogenous Removal of Snickers and Peanut M&Ms

Table 16 reports the detailed product-level results of the joint Snickers-Peanut M&M removal. 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 (i.e., Milky Way and Three Musketeers), but also products from other manufacturers (i.e., Nestle’s Butterfinger). Meanwhile, Raisinets (Nestle), a product that Mark Vend stocks frequently, sees an increase in sales of only 17% when Snickers and Peanut M&Ms are removed, indicating that Raisinets may not be a close competitor to the removed products.⁸⁰ In contrast, 93 consumers substitute to Reese’s Peanut Butter Cups (an 85.6% increase in

⁷⁹Weeks in which the other five treatments were run (for the salty-snack and cookie categories) are excluded from the set of potential control weeks.

⁸⁰Substitution to Raisinets is only 3.3% when Snickers is removed by itself.

sales for the Hershey product), which Mark Vend stocks much less frequently. This provides some descriptive evidence that the rebate may lead Mark Vend to favor products that do not steal business from the major Mars brands over better-selling products that do.

A.6 Results of Nested-Logit Demand Estimation

Table 17 reports the parameter estimates for the nested logit specification, which assumes that $(\mu_{ijt} + \varepsilon_{ijt})$ is distributed generalized extreme value, so that the error terms allow for correlation among products within a pre-specified group.⁸¹ 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:

$$p_{jt}(\delta, \lambda, a_t) = \frac{e^{\delta_{jt}/\lambda_l} (\sum_{k \in g_l \cap a_t} e^{\delta_{kt}/\lambda_l})^{\lambda_l - 1}}{\sum_{\forall l} (\sum_{k \in g_l \cap a_t} e^{\delta_{kt}/\lambda_l})^{\lambda_l}}, \quad (12)$$

where the parameter λ_l governs within-group correlation.⁸² Just as we do for the random-coefficients logit model, we assume $\delta_{jt} = d_j + \xi_t$, and we use five nesting categories: Chocolate, Non-chocolate Candy, Cookie/Pastry, Salty Snack, and Other. Estimation is via maximum likelihood (ML) for the same two definitions of ξ_t used in the random-coefficients specification of table 5.

A.7: Additional Merger Analyses


Table 18 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. Mars acquires Raisinets, which is nearly always stocked by the retailer, and which provides the merged firm with additional profits to leverage. Substitutability between Mars' products and Nestle's products (Butterfinger, Crunch, and Raisinets) is much lower than between the Hershey's product, Reeses Peanut Butter Cups, 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 9, in which the rebate induces the retailer to stock (M,M) instead of (H,H). Table 19 considers the impact of the Hershey-Nestle merger, and achieves

⁸¹See McFadden (1978) and Train (2003).

⁸²Note 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.'

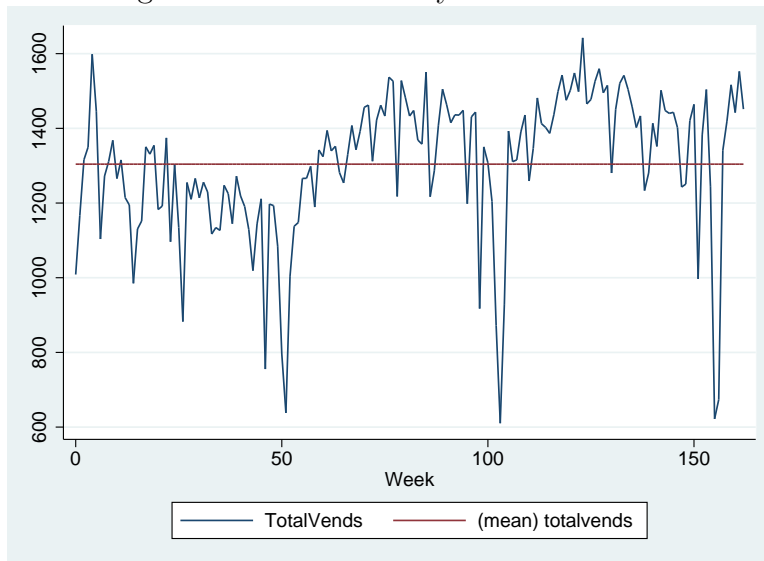
similar results. The only exception is that the merged firm controls more brands and is a somewhat stronger competitor to Mars.

Figure 5: Mars Profits as a Function of Price (Linear Pricing)



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

Figure 6: Overall Weekly Sales at Site 93



Notes: Figures calculated by authors, and represent all product categories in the machines (i.e., confections, snack foods, cookies, and other).

Table 15: Results from Varying Capacity (Numerical 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 M:		
Firm M prefers [1,2]	0.19%	13.78%
Firm M prefers [1,3]	99.55%	85.43%
Mean Firm M profit under [1,2]	17.98	21.52
Mean Firm M profit under [1,3]	20.96	23.62
Firm M pays rebate under [1,2]	0%	0%
Firm M pays rebate under [1,3]	100%	100%
Firm N:		
Mean Firm N profits under [1,2]	8.36	7.73

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

Table 16: Results from Snickers and Peanut M&Ms 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
Reese's 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
Plain M&Ms	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
Raspberry 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	284.5	282.6	-1.9	-0.7	-0.02	-0.10	89
Assorted Baked Chips	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 Chips	198.1	174.6	-23.5	-11.9	-0.29	-1.34	80
Cheeto	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
Peanut M&Ms	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

Notes: Control weeks are defined according to the procedure described in appendix A.4.

Table 17: Nested Logit Estimates

	Parameter Estimates	
$\lambda_{Chocolate}$	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]
# Fixed Effects ξ_t	15,256	2,710
LL	-4,372,147	-4,410,649
BIC	8,972,783	8,862,840
AIC	8,774,962	8,826,873

Notes: The nested logit estimates correspond to the choice probabilities described in section 4, equation 12. Both specifications include 73 product fixed effects. Total sales are 2,960,315.

Table 18: Profits after Mars-Nestle Merger

Policy	Retail (No Reb.)	Rebate	Mars/ Nestle	Hershey	Inte- grated	Industry	Consumer Surplus
Reeses Peanut Butter Cup (H), Three Musketeers (M)							
Retailer-Pre (267)	36,399	2,077	12,978	1,302	49,377	50,679	21,685
Retailer-Post (262)	36,395	2,082	13,013	1,299	49,409	50,708	21,741
Integrated (245)	36,340	2,098	13,114	1,290	49,455	50,745	21,903
Reeses Peanut Butter Cup (H), Payday (H)							
Retailer-Pre (263)	36,661	1,815	11,341	2,173	48,001	50,174	20,845
Retailer-Post (257)	36,656	1,822	11,388	2,167	48,045	50,211	20,914
Integrated (239)	36,591	1,842	11,511	2,151	48,102	50,253	21,086
Three Musketeers (M), Milkyway (M)							
Retailer-Pre (264)	36,090	2,292	14,323	0	50,412	50,412	21,080
Retailer-Post (259)	36,086	2,297	14,354	0	50,441	50,441	21,136
Integrated (244)	36,040	2,310	14,436	0	50,476	50,476	21,281

Notes: Profit numbers represent the long-run expected profit from a single machine in Group D. The two panels represent the two product assortments that offer the greatest potential retailer profit under different stocking policies and rebate payments.

Table 19: Profits after Hershey-Nestle Merger

Policy	Retail (No Reb.)	Rebate	Mars	Hershey/ Nestle	Inte- grated	Industry	Consumer Surplus
Reeses Peanut Butter Cup (H), Three Musketeers (M)							
Retailer-Pre (267)	36,399	1,875	11,719	2,562	48,117	50,679	21,685
Retailer-Post (261)	36,394	1,882	11,763	2,556	48,157	50,713	21,752
Integrated (244)	36,335	1,899	11,871	2,538	48,206	50,744	21,911
Reeses Peanut Butter Cup (H), Payday (H)							
Retailer-Pre (263)	36,661	1,609	10,055	3,458	46,716	50,174	20,845
Retailer-Post (257)	36,656	1,617	10,106	3,449	46,762	50,211	20,914
Integrated (237)	36,578	1,640	10,251	3,421	46,829	50,250	21,102
Three Musketeers (M), Milkyway (M)							
Retailer-Pre (264)	36,090	2,091	13,067	1,256	49,156	50,412	21,080
Retailer-Post (259)	36,086	2,096	13,101	1,254	49,187	50,441	21,136
Integrated (243)	36,035	2,111	13,195	1,246	49,230	50,476	21,289

Notes: Profit numbers represent the long-run expected profit from a single machine in Group D. The two panels represent the two product assortments that offer the greatest potential retailer profit under different stocking policies and rebate payments.