Discussion Paper Series – CRC TR 224

Discussion Paper No. 641
Project B 04

Demand Steering Through the Smokescreen of Stockouts: Evidence from Cigarette Vending Machines

Pablo Casas[1]
Asis Martinez-Jerez[2]
Helena Perrone[3]

February 2025

[1]Arizona State, casas.arce@gmail.com
[2]Cornell University, fm447@cornell.edu
[3]University of Mannheim and MaCCI, helenaperrone@uni-mannheim.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Demand steering through the smokescreen of stockouts: evidence from cigarette vending machines*

Pablo Casas † Asis Martinez-Jerez ‡ Helena Perrone §

November 2024

Abstract

We investigate the use of stockouts as a tool for retailers to steer demand towards higher-margin products. Our empirical evidence suggests that retailers make strategic restocking decisions, putting less effort into restocking low-margin products and prompting consumers to shift purchases towards high-margin products. Our empirical analysis uses a unique dataset where we observe both sales and latent demand, i.e., how many sales a certain product lost when it was out of stock. By exploiting variation in product availability, we recover preference parameters in a setting where prices vary infrequently. Estimated diversion ratios are high across products within the retailer and low towards outside retailers. We also recover manufacturers' marginal costs and perform counterfactual exercises to measure the welfare effects of demand steering on consumers and manufacturers. Our results indicate that while welfare losses are economically relevant on average, some manufacturers benefit from strategic stockouts. Our paper sheds light on the challenges of detecting demand steering in habitual contexts and the market inefficiencies arising from downstream moral hazard.

*Perrone gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG) through the CRC TR 224 (project B04).

†Arizona State, casas.arce@gmail.com

‡Cornell University, fm447@cornell.edu

§University of Mannheim and MaCCI, helenaperrone@uni-mannheim.de

1 Introduction

Product stockouts are prevalent in retail. For example, Gruen et al.(2002) find that around 8% of products sold in stores are out of stock daily in developed countries (see Hickman and Mortimer, 2016, for many other examples in different retail contexts). These stockouts may significantly decrease manufacturers' profits because of the lost sales due to out-of-stock products and consumer surplus if consumers cannot find their favorite products.

Why do these stockouts happen so frequently? Some stockouts may be due to unforeseeable demand shocks, high inventory costs, or low manager effort. This paper proposes and tests a new motive to explain stockouts. We put forward that stockouts may result from profit-maximizing decisions by retailers. By stocking out of lower-margin products, retailers can steer consumers towards higher-margin ones. Stockouts can therefore serve as a non-pricing tool to increase profits. We present empirical evidence of intentional stockouts and quantify their costs for manufacturers and consumers, which can be significant even when the profit gains from demand steering appear small.

Detecting the use of stockouts for demand steering in habitual contexts presents significant empirical challenges. It is difficult to distinguish between unintentional stockouts caused by unpredictable high demand and intentional stockouts aimed at steering demand away from certain products, especially when the stocking efforts of the retailers are unobserved. Additionally, retailers often simultaneously employ other profit-maximizing strategies such as pricing, discounts, promotions, product placing, and changes in the product line, making it difficult to identify the effect of each of them separately.[1] Moreover, conventional vertical agreements between manufacturers and retailers typically used to mitigate conflicts of interest weaken retailers' incentives for demand steering. In these cases, although the incentives for the demand steering mechanism may still be present, they cannot be easily uncovered.

We consider a unique context with features that circumvent most of these challenges and help us identify this strategic motive for stockouts. We study cigarette vending machines in a major European metropolitan area. Each machine sells different cigarette products produced by various multiproduct manufacturers. First, due to the design of our cigarette vending machines, consumers only learn that a product is out of stock after actively trying to purchase the product, by pressing the product button in the machine. Importantly, the machine records that someone pressed the button. Therefore, we observe not only actual sales, but also latent demand for a product when that product is out-of-stock. These data are crucial to identify lost sales and separate stockouts due to high demand and those related to stocking efforts. Second, vending machines are a standard retail format with fixed capacities for a limited number of unique products. Hence, the retailer's decisions regarding assortment and restocking are discrete and relatively straightforward. Furthermore, the assortment is fixed in the medium run, as any change incurs the cost of reprogramming the machine.[2] Third, cigarette prices are set by manufacturers at the national level and vary infrequently (less than once a year). Moreover, regulatory constraints prohibit retailers from engaging in advertisement or promotional activities. Therefore, retailers' primary focus and profit maximization instrument is stocking decisions.[3]

[1] By product line, we mean the products typically offered by the retailer, irrespective of whether they are in or out of stock at a given time. When we mention product assortment, we refer to the products available for purchase at the retailer at a specific time. Hence, it excludes out-of-stock products.

[2] Different from typical snack vending machines, our cigarette vending machines are not see through. The brand available in each channel is indicated by a channel-specific display. Hence, changing brands in cigarette vending machines requires reprogramming each channel to display the correct product. In contrast, see-through snack vending machines do not require reprogramming; product changes are made simply by swapping items in the designated machine channels.

[3] Also, prices do not respond to short-run market-product-specific unobserved demand shocks, simplifying

Fourth, retail markups are fixed, so we observe markups and can study how stocking decisions and stockouts respond to them. Last, due to heavy regulation, vertical agreements between manufacturers and retailers that could alleviate agency conflicts, such as vertical rebates, are prohibited.[4]

We describe a conceptual framework based on the standard newsvendor model. From it, we derive empirically testable implications that allow us to separate the strategic retailer that makes stocking decisions to steer demand towards high-margin products and the naive retailer that overlooks product substitution. In the newsvendor model, a retailer aims to optimize inventory against an uncertain demand. To maximize profits, it must balance the cost of unsold stock when demand is low against potential lost sales when not holding enough inventory to meet demand. In response to this tradeoff, stockouts occur with a positive probability. But because the cost of lost sales increases with the profit margin, the probability of stocking out must be lower for higher margin products.

When dealing with multiple independent products, naive retailers set stock levels independently of their substitution pattern. But strategic retailers can leverage substitution to increase revenues. We can then test whether retailers strategically steer demand by analyzing whether stockout probabilities vary with substitution patterns. We exploit the fact that competition from other retailers decreases product substitution within a retailer, thereby reducing the benefit of strategic stockouts. This is so because a consumer that does not find her most desirable option is less likely to purchase an alternative product in the same machine when they can easily obtain it from an alternative retailer. By analyzing how stockout probabilities relate to retailer competition, we can test if retailers strategically steer demand, with more competition leading to fewer stockouts.

Using product-level daily data from vending machines, we show reduced-form evidence consistent with demand steering through stockouts. First, we show that consumers see different cigarette brands as substitutes. Indeed, diversion ratios from stockout products towards stocked-in products in the same machine are positive and large. Second, we show that the frequency at which a product is out of stock decreases with its margins, controlling for its total demand (sales plus latent demand) and unobserved machine characteristics. Controlling for the product's total demand is essential for this exercise to be meaningful. Otherwise, a negative (positive) correlation between stockout frequency and margins could be due to lower(higher)-margin products having higher demand (and, therefore, running out more frequently).[5] Third, we show that the number of times a product is recharged in a machine per month increases with its margins.

While the evidence so far is also consistent with retailers naively making stocking decision, the next results point more directly to the strategic motive for stockouts. Fourth, monthly machine revenue is higher when there is a stockout, controlling for total sales and machine fixed effects. It is critical to control for total sales in this case: we want to compare two machines with the same number of sales but potentially different sales compositions because one used stockouts to steer demand toward high-margin products, whereas the other did not. Fifth, the probability that a specific product gets recharged increases with margins when at least one product in the same machine gets recharged.[6] Finally, our last result is that stockouts decrease

demand parameter identification.

[4]Although the environment in which we study demand steering is unique, we believe the mechanism is general. Our context is a useful proxy for more complex environments where demand steering is still believed to exist but is harder to detect and quantify. An example is self-preferencing in digital markets. In that example, retailers may hide lower margin products from consumers not necessarily only through strategic stockouts but, e.g., by placing a product on the second page of the search results.

[5]Our empirical evidence indicates that higher margin products have higher demand than lower margin products.

[6]We observe in the data instances when some products in the machine are recharged but not all, including

with competition. This implication is inconsistent with a standard newsvendor model but consistent with demand steering: a retailer can only steer demand from out-of-stock products towards higher-margin products if consumers substitute within the store instead of looking for another store. Therefore, the more store competition there is, the less costly it is for consumers to shop elsewhere, and the less the retailer can use stockouts strategically to increase revenues.

To quantify the welfare costs resulting from stockouts, we estimate preference parameters in a demand model that allows for rich heterogeneity across machines. We take advantage of our panel's long time series and estimate machine-specific parameters combined with a nested logit model where the nests are tobacco type (black, "light," and regular). The demand model explores the observed variation in available products, which is crucial to parameter identification in a context where prices vary infrequently. Using our estimated preference parameters, we construct the counterfactual demands if there are no stockouts. We measure the welfare effects of stockouts by comparing factual and counterfactual indirect utilities. Our estimates show that consumer surplus would be 19% higher in the absence of stockouts.[7]

We then recover manufacturers' marginal costs, assuming they set prices following a multi-product Bertrand model. We combine the marginal costs and estimated preference parameters to calculate manufacturers' profits when there are no stockouts and measure how much profit they would make in this case.[8] We show that the aggregate loss caused by stockouts is relatively small (3%), because most lost sales are diverted to alternative options. But this small aggregate effect hides a large heterogeneity, in which some manufacturers lose up to 21%, while others win up to 14% as a result of the stockouts. The first group are the manufacturers that offer a wide variety of products, while the latter group comprise manufacturers that specialize in the premium segment of the market and, hence, benefit from the strategic actions of the retailers.

A natural question is why retailers resort to strategic stockouts to increase margins instead of adjusting their product assortment permanently. For starters, the assortment of the machine is easier to observe and contract upon, offering a first instrument with which to control the retailer moral hazard problem. But, in addition, we argue that offering only high-margin products would not be profitable for the retailer. In order to attract low-margin product consumers to stores, consumers should not be able to anticipate that low-margin products will be out of stock or not offered at all. If they could perfectly predict it, they would not enter a bar with a machine that never offers their preferred product. Therefore, the retailer must build a reputation for carrying low-margin products with a positive probability (see Dana, 2001; Dana and Petruzzi, 2001; Krishnan and Winter, 2010).[9] Otherwise, the retailer would not benefit from demand diversion from consumers who, once at the bar, prefer to switch cigarette brands over traveling to another bar.

Our paper contributes to the literature on inventory management. There is a vast theoretical work that studies optimal inventory policies. But there is comparably little evidence on how retailers make stocking decisions. A few recent papers show that market structure has an important effect on inventory policies. In particular, Olivares and Cachon (2009) find that car

---

stocked-out products.

[7]The cost of stockouts to consumers disregards health considerations associated with tobacco consumption, treating cigarettes as any other standard non-durable consumer product. We calculate the costs to consumers as the utility costs of consuming a less preferred brand, assuming that stockouts will not impact the quantity consumed and that all cigarette brands have the same health effects.

[8]We consider the short-term or static manufacturers' costs related to immediate substitution for other products. These costs do not include possible costs associated with persistence in choice, where substituting for another product today increases the utility of consuming the alternative brand in the future. This dynamic effect can be especially important in markets like tobacco, where habit formation and brand loyalty are relevant.

[9]A model of assortment decisions, in this case, would therefore require reputation or commitment to carry low-margin products with positive probability.

dealers hold more inventory when competition is stronger. Matsa (2011) shows that supermarkets respond to an increase in competition (due to Walmart opening a store nearby) by reducing stockouts to offer higher quality service to consumers. We extend the literature by showing that strategic stocking decisions, aimed at steering demand toward high-margin products, are a key driver of stockouts. Exploring data on re-stocking decisions, we show direct evidence that at least part of the stockout response to increased competition is due to demand steering motives and cannot be explained uniquely by quality increase.

Smoking out the steering motive behind stockouts is relevant because, in many markets, demand steering can be an antitrust concern due to possible exclusionary effects. Prominent examples relate to demand steering in the form of self-preferencing in digital retailing, such as the recent Amazon Buy Box case. Furthermore, retailers' objectives behind stockouts may not align with the profit-maximizing objectives of the manufacturers. This misalignment is an instance of downstream moral hazard, leading to market inefficiencies that erode manufacturer and consumer welfare.

Our work is close to Conlon and Mortimer (2021b), who study the effect of vertical rebates on retail effort and assortment decisions. Vertical rebates are payments made by the manufacturer to the retailer conditional on some sales target. They are used to relieve downstream moral hazard by incentivizing retailers to exert more sales effort. However, similarly to demand steering, it can have anti-competitive exclusionary effects when, to help attain the target, products from competing manufacturers are inefficiently excluded from the retailer assortment. In their paper, they study a vertical rebate paid by Mars, the leading US candy manufacturer, using data on vending machines selling snacks. They develop and estimate parameters of a model of consumer choice and retailers' dynamic restocking decisions. Identification exploits exogenous assortment variation created by a field experiment where they removed Mars products from vending machines. Empirical results indicate that the vertical rebate led to the foreclosure of competing products and higher profits for Mars but lower consumer welfare and aggregate producer surplus. We complement their evidence by showing that similar exclusionary effects arise non-contractually through strategic stockouts.

Our paper also relates to the literature on moral hazard in expert-client agency contexts. Iizuka (2007) and Iizuka (2012) study the agency problem between doctors and patients in drug prescription using data from Japan, where some doctors prescribe and sell drugs to their patients. The papers examine whether prescription decisions are driven solely by concerns about the patient's welfare or also by drug markups. Results indicate that drug markups significantly affect prescription choices. This evidence is consistent with Liu et al. (2009), who show that financial incentives affect doctors' choice between generics and brand-name drugs in Taiwan, where, like in Japan, doctors prescribe and sell drugs to their patients. However, expert-client relationships are plagued by asymmetric information, making them conducive to demand steering. Our evidence shows that demand steering can be present in markets with strong brand loyalty and in which consumers have strong preferences for purchasing a particular product.

More generally, our work speaks to the current debate on the anti-competitive effects of demand steering and self-preferencing or own-content bias in platforms (see, for example, De Corniere and Taylor (2019); Hagiu and Jullien (2011); Hervas-Drane and Shelegia (2022), among others). Online platforms, for example, can divert demand by posting high-profit margin products (e.g., own products) more prominently or using their recommendation systems to steer consumers towards their own brands. There is growing empirical evidence of steering in online platforms. Farronato et al. (2023); Waldfogel (2024) find evidence consistent with Amazon ranking its own branded products higher than similar products. Similarly, Teng (2022) find that Apple boosts the ranking of its own apps in the App Store. However, the study of online marketplaces poses significant challenges for the identification of demand steering, as we cannot observe counterfactual demand. Because we cannot know what demand would be for lower-ranked products or apps had they been ranked higher, it is difficult to rule out the possibility that consumers prefer the products offered by those platforms, perhaps due to unobserved quality differences. To address this challenge, Teng (2022) develops a structural approach that allows her to estimate the effects of unobserved quality and platform bias, and simulate counterfactual scenarios. In our case, we can do that without imposing as much structure due to the richness of our data and the uniqueness of our context that enables direct observations of the demand for stocked-out products.

The rest of this paper is organized as follows. In the next section, we describe the relevant institutional details of the cigarette market and vending machines in the European metropolitan area we study. Section 3 develops the conceptual framework, and derives the empirically testable implications consistent with a demand steering motive for stockouts. Section 4 describes the data and shows basic statistics, whereas section 5 shows results from the reduced-form tests. Section 6 describes the demand estimation approach, discusses identification and instruments, and shows preference parameter estimation results. In section 7, we describe the counterfactual exercises used to measure the costs of stockouts for manufacturers and consumers. Section 8 concludes.

2 Cigarette vending machines in an European market

We focus on cigarette vending machines in a preeminent metropolitan area of a European country. As with everywhere else in Europe, this country's cigarette market is highly regulated. Tobacco can only be sold in a primary official network of tobacconists who are specially licensed by the state to sell cigarettes. Tobacco can be sold in the secondary network only in vending machines, which are allowed in bars and nightclubs, restaurants, convenience stores, newsstands, and hotels. Tobacco products cannot be sold directly by manufacturers to consumers. In 2017, there were more than 250,000 cigarette vending machines in the country. Around 40% and 50% of total national tobacco sales go through vending machines, which are managed and stocked by tobacconists in the area. Each bar or restaurant that has a cigarette vending machine gets a permit to supply its cigarettes from a specific tobacco store. The bar or restaurant can ask for a different tobacco store but it has to be one of the three closest to it. At any point in time, only one tobacco store is authorized ti supply the vending machine.

Retailers can only stock tobacco in the tobacco store or inside the machine, but never under the counter outside of the machine. Therefore, the maximum stock in bars or restaurants with a cigarette vending machine equals the machine's capacity. There are different machine models, but on average, machines have around 14 channels (i.e., "windows" or buttons for the various brands) and 25 to 35 packs per channel.

Cigarette prices in this country are set at the national level by manufacturers. Any changes in prices must be announced, approved by the government, and published in advance. These changes are infrequent, occurring less than once a year on average. Retailers' margins are also subject to government regulation, fixed at 8.5

In this country, tobacco advertising and promotion are largely prohibited, with a few exceptions. These exceptions allow for limited in-premises actions at tobacco stores, providing a small window for promotional activities in an otherwise heavily regulated market.

3 The naive vs the strategic retailer: conceptual framework

In this section, we use the newsvendor model to develop intuition for how a retailer manages its inventory and optimizes its stock. We also show how we can test for the use of strategic re-stocking to steer demand towards high margin products.

In the newsvendor model, a retailer decides how much inventory to acquire in order to serve customer demand for a single period. The retailer pays a cost of c for each unit of stock acquired, selling it for a price of $p > c$. The units that are not sold at the end of the period are discarded. Demand q follows a continuous distribution $F(q)$, with density $f(q)$. The retailer chooses to stock Q units of product to maximize expected profits, $E\pi = E[p\min(q,Q) - cQ]$, which can be written as:

$$\max_{Q} \int_{0}^{Q} pqf(q)dq + pQ[1 - F(Q)] - cQ,$$

where $1 - F(Q)$ is the probability of stocking out. If we differentiate with respect to Q, the first-order condition yields the optimal inventory level:

$$Q^* = F^{-1}\left(\frac{p-c}{p}\right)$$

The model readily extends to multiple products when there is no substitution. For simplicity, suppose the retailer is stocking two products, $i = 1, 2$, and needs to decide the level of inventory for both, Q_1 and Q_2. Demand for product i follows a distribution $F_i(q_i)$. If the products are not substitutes, demand for a product that stocks out does not convert to sales of the other product. In that case, the retailer can optimize the stock of each product independently, setting inventory levels of $Q_i^* = F_i^{-1}(\frac{p_i - c_i}{p_i})$. When the retailer sets stock levels independently, we say that the retailer is not using stockouts to steer demand.

Notice that the model predicts that the retailer will suffer a stockout of product i with a probability $1 - F(Q_i^*) = c_i/p_i$. Stockouts are the natural response to demand uncertainty when there are costs of holding excess inventory. In the newsvendor model, the cost of excess inventory arises because any unsold units are discarded. At the margin, holding enough inventory to cover demand with probability one is not optimal. The optimal probability of stocking out trades off the cost of holding excess inventory, which equals the cost of acquiring the extra unit, to the cost of lost sales, which equals the margin obtained when selling the additional unit to the customer. When the price (and hence the margin) increases, each lost sale is more costly, and hence, the retailer wants to hold more inventory and reduce the probability of stocking out. The same is true when the product cost is lower, as the cost of holding the excess inventory falls, and the margin increases, making lost sales more costly.

Hence, from this optimal inventory level, we can state the following properties:

Claim 1 *There is a positive probability of stockout,* $1 - F_i(Q_i^*)$*. Moreover, the probability of a stockout decreases in the profit margin; that is, the product with a higher profit margin stocks out less often.*

The model also predicts that revenue is higher when the retailer suffers a stockout than when it does not. This follows from the fact that $E[q_i|q_i > Q_i] > E[q_i] > E[q_i|q_i \leq Q_i]$. That is, sales are higher when a product falls out of stock because stockouts correlate with a higher demand realization.

Claim 2 *Expected revenue conditional on a stockout is higher than expected revenue conditional on no stockout.*

However, sales should not be higher when there is a stockout if we condition on the underlying demand realization q. Indeed, for a given level of demand, revenue should be lower conditional on stocking out than conditional on no stockout.

Suppose now that the two goods are substitutes. We model this by assuming that the retailer sells product i when there is enough stock to meet demand for that product, but a fraction ϕ_i of such customers would purchase product j if available when i is out of stock.

If we keep the retailer's stock levels fixed for both products, it is easy to see that expected revenue increases in each ϕ_i. This is so because, for each realization of demand (q_1, q_2), sales can only increase in ϕ_i: if enough product is available to meet demand, ϕ_i is irrelevant, and if product j stocks out, part of the excess demand is diverted to product i, making sales increasing in ϕ_i. It therefore follows that:

Claim 3 *Holding the stocking decision of the retailer fixed, expected revenue of the retailer increases in the degree of substitution among the two products, ϕ_i.*

However, because substitution only affects sales when one of the products has stocked out, the substitution patterns should not change the probability of stocking out. To see this, notice that the probability that the retailer suffers a stockout equals $Pr(q_1 > Q_1) + Pr(q_2 > Q_{21}|q_1 \leq Q_1)$, which does not depend on ϕ_1 or ϕ_2.

Claim 4 *Holding the stocking decision of the retailer fixed, the probability of having a stockout is independent of the substitution among the two products, ϕ_i, and hence, of the degree of external competition faced by the retailer.*

As a result, we can test for the existence of strategic demand steering by looking at the probability of stocking out. If the retailer is not using stockouts to steer demand and sets stock levels at their individual optimum Q_i^*, such levels should be independent of the degree of substitution ϕ_i. To the extent that the retailer is sophisticated and tries to exploit the consumers' substitution among products, it will alter its stocking decisions based on the likelihood that a consumer that has a preference for a stocked-out product substitutes for another product in stock.

We expect external competition to affect the degree of substitution. When consumers have more alternatives for where to make a purchase, they will be more inclined to search for their preferred product rather than settle for their second-best product at a given retailer. As a result, more external competition for the retailer means a lower product substitution ϕ_i among the retailer's offerings. A sophisticated retailer will respond to an increase in external competition by increasing the level of stock of the products and, hence, lowering the probability of stocking out. We will use this insight to test for the presence of demand steering.

Claim 5 *If the retailer is trying to steer demand, the probability of stocking out will decrease with the degree of competition the retailer faces.*

Although we have developed our intuition using the newsvendor model, which is static in nature, it is important to note that the problem facing retailers is inherently dynamic. They not only choose how much to stock up of each product, but also how often, and when to visit each machine to re-stock them. This is a more intractable problem than the newsvendor model, but our claims should still hold.

In a dynamic re-stocking problem, Scarf et al. (1960) shows it is optimal to follow an Ss policy, restocking when inventory falls below a certain level s, and increasing it to a higher level S. If demand is certain and restocking is instantaneous, $s = 0$ and the retailer never stocks out (Arrow et al., 1951). However, when demand is uncertain and the restocking decision takes place at discrete intervals, it will be optimal for the retailer to tolerate a positive probability of suffering a stockout (Bellman et al., 1955). In that case, the same results as in the newsvendor model should follow: the probability of a stockout should decrease in the profit margin, and the degree of substitution should increase profits, but not change the probability of a stockout, holding the stocking decision of the retailer fixed. Hence, claims 1 to 5 should still hold when the re-stocking dynamics are taken into account.

In addition, if there is a fixed cost of re-stocking the machine, a naive retailer should stock up, not only the products that are out of stock at the moment, but potentially also other products that are likely to stock out before the next re-stocking trip. Indeed, it must be optimal to re-stock at least all products that are out of stock when visiting the machine.[10] Therefore, failing to re-stock those products should be a strong indication of strategic re-stocking. In particular, a strategic retailer can increase the probability of stocking out low margin products by failing to re-stock those products more often than high-margin ones when visiting a machine. Such a retailer may even fail to re-stock a low margin product when it is out of stock, in the hopes of diverting demand towards higher margin ones.

4 Data and descriptive statistics

The data consist of daily information on cigarette vending machines in a large metropolitan area in Europe during 2016, 2017, 2018, and 2019. They include brand, price, other product observable characteristics (e.g., manufacture and pack size), sales, and recharge occasions per machine channel. The data also have information on lost sales, i.e., how many times a consumer tried to purchase a stocked-out product. A peculiarity of these cigarette machines that is very convenient for our analysis is that they are not see-through, unlike regular snack vending machines (see figure 1 for a picture of one of these machines). Therefore, consumers cannot see the products stocked in the machine. They only find out that a product is out of stock after trying to purchase it, i.e., after clicking the product's button. The machine's information system collects data on these frustrated purchases, which we call lost sales. Thus, a unique feature of our data is that we observe the latent demand for out-of-stock products, which other studies can only indirectly estimate. This feature of our data is crucial for separately identifying stockouts resulting from demand shocks and stockouts resulting from re-stocking decisions of the retailer (see section 5).

We also have data on the location of other cigarette vending machines and tobacco stores around each machine. This information will allow us to calculate the degree of competition facing each of the machines from nearby outlets.

For the empirical exercise in this paper, we work with a subsample of 261 machines with the most frequent sales. We consider 15 brands to be "inside" products and bunch the other products in the outside good. The 15 inside products have a joint market share of close to 90%. The manufacturers of the inside products are the four leading tobacco manufacturers in the world market: Imperial Tobacco, British American Tobacco, Phillip Morris, and Japan Tobacco.

Some machines have multiple channels for the same product. Hence, one channel may be out of stock (and the machine records lost sales for that channel), whereas the product is still available in another channel of the same machine.[11] For that reason, we further restrict our subsample by eliminating those machine-day combinations when a product was out of stock in one specific channel but not in the machine as a whole. To check for robustness, we also performed the reduced-form tests using the larger subsample, and the results remain qualitatively the same.

Our final subsample has 2,180,645 observations; one observation being a product-machine-day combination (we aggregate different channels selling the same product). These observations correspond to 119,919 machine-day combinations in the 261 distinct machines. Table 1 shows

[10]Otherwise, if it was optimal to leave a product out of stock when paying the fixed cost of visiting the machine, the retailer would be better off not carrying that product, and instead, devoting the machine space to a different one.

[11]When a cigarette machine has multiple channels selling the same product, the consumer must click on each channel to find out if it is in stock. Clicking on an out-of-stock channel does not transfer a product from a stocked channel to an out-of-stock one.

summary statistics of the variables relevant to our study at the machine-day level. The mean number of distinct products sold in a machine is 19, and the mean total daily sales per machine is slightly above 16 packs, thus less than one pack per product. The mean daily revenue per machine is 78 euros, so the average price per pack sold is 4.8 euros. On any day, the probability that a machine has at least one product out of stock is around 57%. This probability is calculated as the average, across machines, days, and products, of a variable equal to 1 if a product is out of stock and zero otherwise. The probability of recharge, calculated as the average across machines and days of a variable equal to one if at least one product in the machine is recharged and zero otherwise, is around 18%, much lower than the probability of a stockout. Table 1 also shows summary statistics on the number of days between recharges (12 on average), lost sales at any given day (4 on average), number of products that stockout on any given day (1.2 on average), and on the density of machines and of tobacco stores in a radius of 500 meters around each machine (27 machines and 3,5 tobacco stores on average).

Table 2 shows summary statistics for each of the 15 inside products we consider in the empirical analysis. It includes market shares, mean prices, mean sales, probabilities of stockouts and recharges, and metrics related to product recharges. Market shares range from 2% to 15%, and the total market share held by the 15 inside products is 86%. Mean prices vary from 4.25 to 5.10, and mean sales per day range from 0.45 to 2.62 units. The probability of a stockout for a given product on any day ranges from 1 to 85%. Metrics between recharges include the average number of days (16.35 to 118.18), number of sales (8.01 to 23.90), and revenue (34.08 to 107.79) between recharges for each product.

5 Reduced-form evidence consistent with strategic stockouts

This section shows empirical evidence consistent with strategic stockouts. We show that retailers use stockouts to increase revenue when different brands are substitutes. We also show empirical evidence that helps reject the alternative hypothesis that observed stockouts result from retailers making stocking decisions taking the demand for each product independently, without leveraging the substitution across products.

We start by showing that there is substitution across brands. We provide nonparametric evidence that diversion ratios between products are greater than zero and sizeable. However, it is possible that, even though consumers see different cigarette brands as substitutes, retailers overlook this and make stocking decisions for each product independently. Thus, this first piece of evidence is necessary but not sufficient to demonstrate the demand steering stockout story. Furthermore, we show that: (i) machine revenue is higher when there are stockouts, controlling for total sales; (ii) stockout probabilities decrease with margins, controlling for total demand; (iii) retailers facing stronger local competition stockout less frequently; and (iv) out-of-stock products with higher margins are recharged more frequently, controlling for total demand. The last two results are critical to distinguish demand steering stockout motives from standard newsvendor motives because competition does not affect stocking decisions when the retailer ignores the substitution across brands, but it does when the retailer makes stocking decisions strategically to steer demand to certain products. In addition, a naive retailer ought to re-stock all out-of-stock products when visiting a machine, regardless of their margin.

Remark that the validity of our demand steering tests depends crucially on us observing total demand, that is, the sum of observed sales and the number of lost sales of out-of-stock products. These data permit separating stockouts resulting from high demand from stockouts related to strategic stocking decisions.

5.1 Structure-free diversion ratios

In this section, we estimate diversion ratios without imposing structure, exploiting the observed variations in product availability resulting from stockouts. These diversion ratios are akin to measures calculated using survey data where respondents indicate which products they would switch to if their preferred product exited the market.[12] Diversion ratios measured this way, i.e., as resulting from an infinite increase in the price, may differ from diversion ratios resulting from marginal changes in price. However, they still capture substitutability across products, which is our main objective in this exercise.[13]

Table 3 shows pairwise diversion ratios averaged across machines. The main diagonal of the table is always zero because it measures the diversion of a product to itself so, by definition, it is zero. Other diversion ratios are typically above zero. The estimated aggregate diversion ratios toward other inside brands ranged between 50% and 100%, indicating significant substitution between different products. Exceptions to this are products 3 and 13, with diversion ratios from these brands to other products totaling less than 40%. Product 3, in particular, exhibits the lowest substitution toward other brands, which is unsurprising given that it is the only cigarette product among the inside brands made from black tobacco. This differentiation likely makes it less substitutable compared to products made from blond tobacco.

5.2 Machine revenues and stockouts

We argue that stockouts result from profit-maximizing strategies of retailers who allocate less effort to avoid stockouts of low-margin products than of high-margin products. This strategy can be profitable to retailers when demand for the out-of-stock product diverts towards higher-margin products. In this section, we show that incentives for such strategic behavior exist.

Suppose this form of strategic stockout is indeed happening in our data. Compare two identical machines with the same sales in a certain period, but one has stockouts, and the other does not. Then, although the number of product sales is the same between the two machines, the composition of the sales differs because, on average, the machine using strategic stockout sells a higher share of high-margin products. Hence, machines that experienced stockouts should have higher revenues than machines that did not. This correlation between stockouts and revenue also happens in the newsvendor model. Therefore, the empirical evidence in this subsection only shows that incentives exist for retailers to steer demand towards high-margin products through stocking out of low-margin ones. It is necessary evidence but not sufficient to distinguish the strategic retailer from the naive one.

To test this possibility, we run regressions of machine per period revenue on whether there was a stockout in the period (month). Critically, we control for sales, machine unobservable characteristics (machine fixed effects), and common unobservable factors (year and month fixed effects).[14] Table 4 shows the results from these regressions. The first column shows estimated coefficients from regressing monthly machine revenues on whether there was a stockout that month without further controls. In column 2, the regression includes machine fixed effects.

[12] This type of survey is used by competition authorities in Europe, e.g., the UK's CMA.

[13] See Conlon and Mortimer (2021a) for a comprehensive discussion on the properties of diversion ratio and under which conditions diversion ratios measured with infinite price variation approximate diversion ratios following marginal price variations.

[14] Remark that when we are looking at the machine level, we should look at total sales, not sales plus lost sales because as there is substitution from lost sales to other brands within the machine, the sum of lost and actual sales will overestimate actual demand at the machine level. When the analysis is at the product level, the relevant measure of demand is sale plus lost sales because this is the variable that measures the total number of consumers that wished to purchase that product in that machine that day (and some could not purchase it when the product was out of stock).

Column 3 also controls for monthly machine sales and includes year fixed effects.

The coefficient for stockout in column 1 is positive and significant. The coefficient sign in column 1 is as expected because stockouts increase with sales. Hence, unobserved positive demand shocks raise revenues and stockout incidence. Thus, controlling for total sales and including fixed effects to control for the remaining unobserved shocks is crucial for identifying the effects of strategic stockouts on revenues. When we include further controls (sales, machine, and month and year fixed effects) and the coefficient estimates are identified from comparable machines with the same level of demand, the magnitude of the estimated stockout coefficients decreases but remains positive and significant. These results imply that, all else equal, machines that have stockouts have higher revenue.

The estimated coefficient for stockout in the last column is small, but we should be careful when interpreting it. This estimate is only a lower-bound for the benefits of strategic stockouts, as re-stocking the machine is costly, both in terms of the transportation cost to travel to the machine as well as the cost of holding the inventory in it. These costs can be lowered when re-stocking less frequently.

5.3 Stockout probability and margins

Suppose retailers indeed make a lower effort to restock lower-margin products. Then, lower-margin products should stock out more frequently, controlling for product total demand (product sales plus lost sales). It is crucial to control for total demand, not only sales. Otherwise, a negative correlation between margins and stockouts could be due to lower-margin products having higher demand, for example, even if restocking decisions are independent of retail margins.[15] This negative correlation between stockouts and product margins also comes up in the naive newsvendor model, where stockout decisions ignore product substitution. Therefore, empirical evidence pointing to higher stockout frequency of lower margin products is necessary but not sufficient for our story of strategic stockout.

Table 5 shows estimation results of linear regressions of product stockout on product margin (results from probit are qualitatively the same). The dependent variable y_{jmt} in each regression is a discrete variable equal to 1 if product j is out of stock in day t and machine m, and zero otherwise. The main right-hand side variable is product j's margin (in euros), corresponding to 8.5% of the product's price. The first column shows the results of a regression of the stockout variable on margins without any other control. Columns 2, 3, and 4 include controls for the product's total demand in day t and machine m. Columns 3 and 4 also include controls for whether the day is a weekend and manufacturer fixed effects. Column 4 adds machine fixed effects.

When we do not control for total demand (column 1), the estimated coefficient for the margin is positive and significant, indicating a positive correlation between margins and stockouts. However, once we control for total demand (columns 2, 3, and 4), the margin coefficient becomes negative and significant. This result implies that, if we compare two products with the same demand, the product with lower margins will stock out more frequently, consistent with retailers exerting higher effort to avoid or reduce stockouts of higher-margin products.

5.4 Competition and stockouts

In this section, we present evidence that is consistent with retailers being strategic but inconsistent with retailers naively ignoring the substitution that consumers make when a product is

[15] Lower-margin products have lower prices than higher-margin products. Hence, a negatively sloped demand would explain higher stockout rates.

out of stock.

To increase profits through strategic stockout of low margin products, retailers need some degree of market power. Otherwise, consumers faced by a stockout of their first choice, will not substitute within the machine but look for another retailer. This insight implies that we should observe more stockouts in machines with lower competition from other cigarette retailers, all else equal. More geographically isolated machines (away from other machines or tobacco stores) can benefit more from stockouts because the transportation costs for consumers to look for alternative retailers are higher than when there are many machines around. Empirical evidence linking stockouts and competition is crucial to distinguish between the strategic retailer and the naive newsvendor, as discussed in detail in section 3.

We study, therefore, how stockout probabilities relate to the density of competitors at a radius of 500 meters around the machine. Table 6 shows the results of regressions where the dependent variable is equal to one if there was a stockout in a machine and a specific month. The main right-hand-side variable is the number of tobacco stores in a radius of 500m around machine m. Our main measure of competition is the number of tobacco stores in the area because we have data on the population of tobacco stores. In contrast, we only have the location of a sample of machines (cigarette machines of the same brand as ours). Hence, the number of machines in the area does not measure retail competition well (we nevertheless include this variable as a control). Column 1 shows that the density of tobacco stores negatively correlates with the monthly stockout frequency, consistent with our demand steering theory.

In column 2, we add a control for total sales. It is essential to do so because a high number of machines and tobacco stores in an area can be due to high demand for cigarettes in that area. Failing to control for sales may bias our inference if, for instance, high-demand areas have a higher number of cigarette machines and stores, leading to a positive correlation between the number of competitors and stockouts. Indeed, when we control for demand, our estimate of the effect of the density of tobacco stores becomes more negative.

Nevertheless, proximity to tobacco stores could also decrease stockouts because it may decrease the cost of recharging the machine (perhaps because the retailer responsible for recharging it is closer). To control for the possibility that the negative correlation between stockouts and the number of retailers could be solely due to retailers recharging nearby machines more frequently, we control for the recharge frequency. Once we do that (column 3), the estimated coefficient for tobacco stores' density remains negative and significant, indicating that lower recharge costs cannot explain the estimated negative correlations. Indeed, the coefficient estimate for the density of tobacco stores remains very similar to the previous specification.

5.4.1 Recharge and margins

In the previous exercise, we used stockout probabilities controlling for total demand to examine the recharge effort of the retailer indirectly. Now, we look at recharge opportunities directly. There are two ways retailers could neglect effort to avoid low-margin product stockouts. First, retailers could wait longer to visit a machine to restock a low-margin product. Second, retailers could refrain from recharging a stocked-out low-margin product even when recharging other products in the machine. In this section we present more direct evidence that points to the strategic withholding of re-stocking effort for low-margin products by retailers.

We start by studying whether the time interval between two product recharges decreases with the product's margin. To do that, it is critical to control for total demand because, even if recharges do not respond to margins, we still expect products with a higher demand to be recharged more frequently. This is so for two reasons. First, the stock of available product depletes faster when demand is high. In addition, keeping a product out of stock is more costly when demand is high, as the retailer suffers higher lost sales. Controlling for total demand therefore is necessary to make sure we capture the retailer's response to margins, rather than the underlying demand.

Table 7 shows estimated coefficients for a regression of the number of days between product j's recharges in machine m and its margins. The coefficient in the first column indicates that the duration between recharges decreases with margins. However, this could be due, for example, to higher margin products having higher demand, requiring more frequent recharges. To check whether this is the case, columns 2 and 3 control for product total demand (sales plus frustrated sales). The magnitude of the coefficients is reduced slightly, suggesting a positive correlation between demand and margins. However, the negative correlation between the recharge interval and margins remains once we control for total product demand (column 2) and even when we add machine fixed effects (column 3).

Next, we check whether the probability that a stocked-out product is recharged increases with margins, conditional on the machine getting recharged (that is, there is another product that is re-stocked in that machine that same day). Table 8 shows results for a linear probability regression of a product getting recharged (the dependent variable is equal to 1 if product j is recharged at period t in machine m) and the product's margin, in which we restrict the sample to those machine-days with at least one product being recharged. The coefficient estimate in column 1 shows a positive and significant correlation between margins and a product's recharging probability. This result indicates that, even when the retailer visits a machine to recharge it, she does not necessarily recharge all products, especially not those with low margins. Columns 2 and 3 constrain the subsample further to include only out-of-stock products, with column 3 adding machine-fixed effects. The coefficients remain positive and significant. This is notable, as it shows that retailers refrain from recharging low-margin products even when the low-margin product is out of stock at the time that the recharging visit occurs. Because the decision to travel to a machine to restock likely entails a sizeable fixed cost, we would expect the retailer to, at least, replenish all out-of-stock products. The fact that they don't do so, particularly for low-margin products, is a strong indication that retailers are being strategic.

Overall, the results suggest that retailers use stockouts strategically. Although some of our non-parametric evidence is consistent with naïve retailers ignoring demand substitution, our latest results are not. They require a strategic motive for stockouts aimed at steering demand towards higher-margin products. Next, we develop and estimate a flexible structural model to quantify the welfare implications of such strategic stockouts.

6 The welfare costs of stockouts

In this section, we quantify the welfare costs for manufactures and consumers resulting from product stockouts. To do that, we develop a nested demand model for cigarette products and a multiproduct-Bertrand model of manufacturer supply. Using the demand model, we estimate consumers preference parameters and simulate product-level demand under the counterfactual scenario of no stockouts. Comparison of the associated factual and counterfactual consumer surplus allows measurement of the surplus costs for consumers of having to substitute from their favorite products that are out of stock to in-stock alternatives. The preference parameters are also used in the multi-product Bertrand first-order conditions to recover manufacturers marginal costs and to calculate counterfactual firm-specific profits in the counterfactual scenario of no stockouts. In this way, we can calculate the profit loss (or gain) associated with products being out of stock in the cigarette machines.

6.1 Demand

We start by estimating a flexible demand model to recover consumer preference parameters. Our demand model takes advantage of the long temporal dimension in our panel to estimate machine-specific preference parameters. Our approach allows for full unobserved consumer heterogeneity across machines combined with a 1-level-nested logit, where groups are defined by tobacco type (regular, "light', and black). Compared to the standard random coefficient approach, we model consumer preferences as individual level parameters that we estimate, rather than treating them as random effects drawn from a known distribution.[16] We also allow for consumption sets to vary per machine and period (due to stockouts and assortment changes), which represents an essential source of data variation for parameter identification in a setting where prices vary infrequently.

We define the indirect utility of a consumer as:

$$u_{ijt} = \alpha_i p_{jt} + \xi_{ijt'} + \gamma_t + \tilde{\epsilon}_{ijt} \tag{1}$$

where p_{jt} is the price of product j at period t, α_i is consumer i's marginal utility of income, ξ_{ijt} is consumer i's taste for unobserved product characteristics of product j that may vary over time, γ_t captures common factors, and $\tilde{\epsilon}_{ijt}$ are consumer- and product-specific per-period unobserved shocks.

We make the nested logit model distributional assumption on $\tilde{\epsilon}_{ijt}$, which allow consumer product valuations to be correlated among products in a same group (See Verboven, 1996, for a detailed discussion of the nested logit model). We consider three groups depending on the type of tobacco (regular, "light', and black), plus the outside good. We assume that

$$\tilde{\epsilon}_{ijt} = \zeta_{igt} + (1 - \sigma_i)\epsilon_{ijt} \tag{2}$$

where ϵ_{ijt} is iid extreme value and ζ_{igt}, where g indexes the group, has a distribution such that $\tilde{\epsilon}_{ijt}$ is extreme value. The nesting parameter σ_i measures consumer i's taste correlation across products in the same group and its value should lie between zero and 1. Notice that the nesting parameter may be individual-specific. Hence, our demand model allows for a rich pattern of consumer heterogeneity.

In period t, consumer i chooses to purchase the product that maximizes her indirect utility. Hence at every period t, the probability that i chooses j is

$$s_{ijt} = P(y_{ijt} = 1 | \alpha_i, \gamma_t, \boldsymbol{\xi_{it}}, \boldsymbol{p_t}) = P(u_{ijt} \geq u_{ikt}, \forall k \in A_{it} | \alpha_i, \gamma_t, \boldsymbol{\xi_{it}}, \boldsymbol{p_t}) \tag{3}$$

where $y_{ijt} = 1$ if consumer i purchases j at period t and zero otherwise, and $\boldsymbol{\xi_{it}}$ and $\boldsymbol{p_t}$ are vectors stacking the unobserved taste parameters and product prices, respectively. Given the above distributional assumptions and normalizing the mean utility of the outside good to zero, this probability can be written as (Ivaldi and Verboven, 2005):

$$s_{ijt} = \frac{\exp(\delta_{ijt})/(1 - \sigma_i)}{D_{igt}} \frac{\exp D_{igt}^{(1-\sigma_i)}}{1 + \sum_{g=1}^{G} \exp D_{igt}^{(1-\sigma_i)}} \tag{4}$$

where $\delta_{ijt} = \alpha_i p_{jt} + \xi_{ijt'} + \gamma_t$ and

$$D_{igt} = \sum_{k \in G_g} [\exp \delta_{ijt}/(1 - \sigma_i)] \tag{5}$$

[16] See Grigolon and Verboven (2014) for a thorough comparison between the nested logit and random coefficients models

where G_g is the set of products in group g.

We assume ϵ_{ijt} in equation (1) has the one-level "nested logit" distribution. Aggregating individual probabilities at machine-level and doing the standard inversion (Berry, 1994), we get the estimable equation:

$$\ln(s_{mjt}/s_{m0t}) = \alpha_m p_{jt} + \sigma_m \ln(s_{jmt|g_m}) + \xi_{mjt'} + \gamma_m + \epsilon_{mjt} \quad (6)$$

where m indexes machine, $s_{jmt|g_m}$ is the market share of product j within its group in machine m and period t, and σ_m captures machine-specific unobserved taste correlations between products within the same group.

6.2 Empirical implementation and preference parameter estimates

In general, prices are endogenous in a demand equation. Firms' pricing decisions respond to unobserved product characteristics and period-specific unobserved shocks to the willingness to pay. Therefore, they are typically correlated with the error term in the demand equation. In equation (6), we explicitly control for products' machine-specific unobserved characteristics and allow them to vary over time (ξ_{mjt}). Conditional on this rich set of controls for consumers' product taste heterogeneity, prices can be considered exogenous in our setting. That is because, as mentioned earlier, cigarette prices in the country we study are set by the manufacturers at the national level and vary infrequently. Therefore, they do not respond to short-term machine-specific shocks to consumers' willingness to pay for the product. Thus, they are uncorrelated with ϵ_{ijt} conditional on including product fixed effects that capture unobserved product characteristics.

However, we should consider the endogeneity of product j's market share within its group, $s_{jmt|g_m}$. This share can be correlated with unobserved shocks affecting the probability of choosing product j at a certain period t and machine m. Therefore, unbiased and consistent estimates of model parameters require using instrumental variables correlated with the group shares but uncorrelated with the willingness to pay for product j. We follow the standard strategy in the literature (Ivaldi and Verboven, 2005) and exploit group-level variation in available sets.

In our application, we should also discuss another potential endogeneity problem. We argue that product stockouts are partly due to strategic profit-maximizing retailers' decisions. Hence, available products could be correlated with unobserved per-period and per-product demand shocks in some periods and machines. Assume retailers' stocking decisions are affected by a per-period and per-machine overall unobserved shock, not a product-specific shock. An example of such a shock would be the arrival of consumers with high transportation costs unwilling to switch machines if their favorite product is out of stock. Then, including day fixed effects that control for unobserved common factors at the machine-day level solves the issue.

Remark that our strategic stockout story says that retailers' stocking decisions are affected by products' retail margins and prices. Therefore, there is a correlation between available products and prices in some periods. However, as we explicitly control for prices, this correlation does not challenge identification (it would be different if prices responded to unobserved period- and product-specific shocks to demand. In this case, some unobservable shock to demand could affect both the set of available products and the prices, creating an endogeneity problem.)

Another potential threat to identification is the following. Suppose a retailer offers a product set Ω when she expects a selected group of consumers to show up at a particular day t in her machine (or the same consumers to behave differently in t). Hence, the group of consumers on day t could be different in unobservable ways to consumers that show up when Ω' is offered. In this case, the preference parameters in t differ from those in t'. Our identification assumption is that conditional on machine-specific and date-specific unobserved shocks, consumers' preferences are comparable across periods in the same machine.

We estimate a demand equation for each machine. Hence, we estimate machine-specific taste parameters that should reflect taste differences between customers who shop in different machines. We include prices for each product, product fixed effects to capture unobserved product characteristics, and product-year and product-month fixed effects, allowing product perception to evolve over time and seasonally. We also include day-specific fixed effects. We use the number of available products within each group, machine, and day as the instrument for group market shares.

Table 9 shows means and standard deviation values for the estimated preference parameters α_m, the price coefficients, and σ_m, the group market share coefficient. Both means are within the expected intervals for the parameters to be consistent with utility maximization, i.e., $\alpha_m < 0$ and $\sigma_m \in [0, 1]$. Standard deviations around these mean values are relatively high, suggesting relevant heterogeneity across machines concerning price sensitivity and taste correlations within a product group.

6.3 Effect of stockouts on consumer surplus

In this section, we use our preference parameter estimates to construct the market shares under the counterfactual scenario where there were no stockouts. In that way, we can calculate the consumer surplus in this counterfactual scenario and compare it with the factual consumer surplus, measuring the costs of stockouts to consumers.

Although we observe in the data the number of frustrated purchases of an out-of-stock product, we do not observe where these frustrated purchases went (which other product, if any, in the machine). Therefore, we do not directly observe what per-product sales would have been without stockouts. For that, we need to use our demand model's estimated consumer preference parameters to simulate consumers' choices when there are no stockouts (conditional on the machine product line). We can then calculate consumer surplus with stockouts (observed choices) and when there are no stockouts (simulated choices using estimated preference parameters) to measure the consumer welfare costs of stockouts.

Given our demand model assumption, the net consumer surplus, CS, is measured by:[17]

$$CS = \frac{1}{\alpha} ln \left(1 + \sum_{g=1}^{G} D_g^{1-\sigma} \right) \tag{7}$$

Table 10 shows results on the effect of stockouts on consumer surplus. Without stockouts, consumer surplus would have been 19% higher on average.

6.4 Effect of stockouts on manufacturers profits

We now turn to the manufacturers to estimate how the re-stocking effort of the retailers affects their welfare. We combine the counterfactual market shares calculated in the previous subsection with assumptions on manufacturer conduct to recover manufacturers variable profits when there are no stockouts. We then calculate the difference in variable profits under the factual and counterfactual scenarios.

We assume manufacturers play a static multiproduct Bertrand game and that marginal costs are constant. Assume also that the cigarette demand at the machines is representative of the national cigarette demand in general (because prices are set nationally). Then, each manufacturer f set prices of each of the products $j \in F_f$ following first order condition:

[17] See Ivaldi and Verboven (2005).

$$\sum_m M_m s_j^m(p) + \sum_{k \in F_f} (p_k - c_k) \sum_m M_m \frac{\partial s_k^m(p)}{\partial p_j} = 0 \tag{8}$$

The set of J first-order conditions allow us to solve for each of the J products marginal costs explicitly. Define $S_{jk} = -\sum_m M_m \frac{\partial s_k^m(p)}{\partial p_j}, j, k = 1, ..., J$,

$$\Gamma^\star_{jk} = \begin{cases} 1, & \text{if } \exists f : (k, j) \subset F_f, \\ 0, & \text{otherwise} \end{cases}$$

and Γ_{jk} is a JXJ matrix with $\Gamma_{jk} = \Gamma^\star_{jk} * S_{jk}$. In vector notation, the marginal costs are:

$$c = p + \Gamma_{jk}^{-1} s(p) \tag{9}$$

where $s(p)$ is a $JX1$ vector such that $s(p)_j = \sum_m M_m s_j^m(p)$, and p and c are $JX1$ vectors of prices and marginal costs, respectively.

After recovering the marginal costs, we can then estimate the impact of stockouts on the profits of each of the manufacturers. Results are in the second panel of Table 10. We find that on average (across manufacturers), the loss from stockouts is relatively small, around 1%. The total manufacture profit loss associated with retail stockouts (i.e., the sum of profit losses in a year for all manufacturers) is 3%.[18] However, these numbers hide relevant heterogeneity across manufacturers. Because demand is diverted from out-of-stock products towards other available products from competing manufacturers, some firms lose from stockouts, while some benefit from them. We find that manufacturers 1 and 2 would be better off in the counterfactual scenario of no stockouts (obtaining 9% and 21% higher profits in the counterfactual, respectively). In contrast, firms 3 and 4 benefit from stockouts, and their profits would have been 14% and 2% lower, respectively, if there were no stockouts. These results are consistent with the fact that manufacturers 1 and 2 offer a wider variety of products at different price points, some of them cheaper products, whereas manufacturers 3 and 4 are more specialized in "premium," more expensive products. Therefore, 1 and 2 lose more from retailers' strategic stockouts aiming at shifting demand from cheaper to more expensive products.

Notice that the counterfactual exercise is static. Hence, the effects of stockouts on manufacturer surplus only measure short-term losses. The dynamics effects of stockouts could be different. For example, stockouts could affect manufacturers' product line choices and prices. They could also accentuate certain manufacturers' profit losses if the demand steering they induce affects brand loyalty and, therefore, long-term product demand, eventually leading to the foreclosure of certain brands. These considerations fall beyond the scope of our study.

7 Conclusion

Using data from cigarette sales in vending machines, we show empirical evidence consistent with demand steering through strategic stockout. Retailers benefit from stockouts by changing product assortment in the short run and diverting demand from out-of-stock products to products with higher retail margins.

The cigarette market is highly regulated, preventing the use of vertical agreements to solve agency problems. Therefore, it constitutes a unique setting for studying the extent and the

[18]Notice that the average loss weights the loss of all manufacturers equally, whereas the total manufacturer profit loss takes into consideration each manufacturer's market share. The latter is larger than the former because, as described later, manufacturers that produce lower margin products have higher market shares and suffer larger losses.

costs of downstream moral hazard for manufacturers and consumers. Furthermore, we observe not only actual sales but also frustrated demand for out-of-stock products, which is crucial to separate stockouts due to unexpected demand shocks from retailers' re-stocking decisions without imposing further structure. Our setting is also unique because we observe frequent variations in consumption sets due to frequent stockouts, which enables the identification of preference parameters in an industry with no price variation. Our counterfactual exercises indicate that the costs of stockouts to manufacturers and consumers are sizable, even when the retailer profit advantage of using demand steering is small. Moreover, although the overall effect of stockouts on manufacturers' profits is negative, some manufacturers benefit from them. These manufacturers offer more expensive products, so their products are likely targets of the retailers' demand steering strategy.

Our paper and results have relevant public policy and managerial implications. Regarding public policy, competition practitioners should consider that frequent stockouts could be a smokescreen for demand steering with potentially exclusionary consequences. For management, our research underscores the importance of considering retailer strategic stockouts when setting manufacturer prices. It also sheds light on an alternative non-price profit maximization tool for retailers, especially when the product line is fixed in the short run.

References

Arrow, K. J., T. Harris, and J. Marschak (1951). Optimal inventory policy. *Econometrica: Journal of the Econometric Society*, 250–272.

Bellman, R., I. Glicksberg, and O. Gross (1955). On the optimal inventory equation. *Management Science 2*(1), 83–104.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 242–262.

Conlon, C. and J. H. Mortimer (2021a). Empirical properties of diversion ratios. *RAND Journal of Economics 52*(4), 693–726.

Conlon, C. T. and J. H. Mortimer (2021b). Efficiency and foreclosure effects of vertical rebates: Empirical evidence. *Journal of Political Economy 129*(12), 3357–3404.

Dana, J. D. (2001). Competition in price and availability when availability is unobservable. *RAND Journal of Economics*, 497–513.

Dana, J. D. and N. C. Petruzzi (2001). Note: The newsvendor model with endogenous demand. *Management Science 47*(11), 1488–1497.

De Corniere, A. and G. Taylor (2019). A model of biased intermediation. *Rand 50*(4), 854–882.

Farronato, C., A. Fradkin, and A. MacKay (2023). Self-preferencing at amazon: Evidence from search rankings. In *AEA Papers and Proceedings*, Volume 113, pp. 239–43.

Grigolon, L. and F. Verboven (2014). Nested logit or random coefficients logit? a comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics 96*(5), 916–935.

Hagiu, A. and B. Jullien (2011). Why do intermediaries divert search? *Rand 42*(2), 337–362.

Hervas-Drane, A. and S. Shelegia (2022). Retailer-led marketplaces.

Hickman, W. and J. H. Mortimer (2016). Demand estimation with availability variation. In *Handbook on the Economics of Retailing and Distribution*, pp. 306–340. Edward Elgar Publishing.

Iizuka, T. (2007). Experts' agency problems: evidence from the prescription drug market in japan. *The Rand journal of economics 38*(3), 844–862.

Iizuka, T. (2012). Physician agency and adoption of generic pharmaceuticals. *American Economic Review 102*(6), 2826–58.

Ivaldi, M. and F. Verboven (2005). Quantifying the effects from horizontal mergers in european competition policy. *International Journal of Industrial Organization 23*(9-10), 669–691.

Krishnan, H. and R. A. Winter (2010). Inventory dynamics and supply chain coordination. *Management Science 56*(1), 141–147.

Liu, Y.-M., Y.-H. K. Yang, and C.-R. Hsieh (2009). Financial incentives and physicians' prescription decisions on the choice between brand-name and generic drugs: Evidence from taiwan. *Journal of health economics 28*(2), 341–349.

Matsa, D. A. (2011). Competition and product quality in the supermarket industry. *The Quarterly Journal of Economics 126*(3), 1539–1591.

Olivares, M. and G. P. Cachon (2009). Competing retailers and inventory: An empirical investigation of general motors' dealerships in isolated us markets. *Management science 55*(9), 1586–1604.

Scarf, H., K. Arrow, S. Karlin, and P. Suppes (1960). The optimality of (s, s) policies in the dynamic inventory problem. *Optimal pricing, inflation, and the cost of price adjustment*, 49–56.

Teng, X. (2022). Self-preferencing, quality provision, and welfare in mobile application markets.

Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics*, 240–268.

Waldfogel, J. (2024). Amazon self-preferencing in the shadow of the digital markets act. Technical report, National Bureau of Economic Research.

Figure 1: Example of cigarette vending machine

Table 1: Summary Statistics at Machine x Day level

	mean	sd	min	max
Number of machines	261	0	261	261
Number of brands x machine	18.902	4.735	1	37
Daily sales x machine	16.558	16.043	0	249
Daily revenue x machine (€)	78.718	75.954	0	1191
Daily lost sales x machine	3.939	8.018	0	117
Daily stockout x machine	1.282	1.678	0	15
At least one product stocked-out	0.573	0.495	0	1
Probability of machine recharge	0.180	0.384	0	1
Number of days between machine recharges	8.233	12.540	1	120
Density of machines in 500m radius	27.310	12.532	1	78
Density of tobacco store in 500m radius	3.459	1.997	1	20
N	127104			

One observation is one machine-day; Density of machines/tobacco stores in radius equals he number of other machines/tobacco stores in 500 m radius around each machine.

Table 2: "Inside" products' summary statistics

Products	Market Shares	Mean Price	Mean Sales	Prob of stockout	Mean lost sales	Prob of recharge	Sum recharge	Between recharges nb of days	Between recharges nb of sales	Between recharges revenue
1	0.13	4.80	2.19	0.18	0.72	0.15	2.27	18.51	16.37	78.62
2	0.02	4.80	0.61	0.08	0.17	0.08	0.85	62.90	9.56	46.02
3	0.02	4.79	0.45	0.04	0.06	0.07	0.67	118.18	9.04	43.38
4	0.08	5.00	1.69	0.21	0.85	0.14	1.89	28.55	11.88	59.44
5	0.04	4.55	0.83	0.07	0.17	0.10	1.01	48.09	9.22	42.00
6	0.07	4.70	1.14	0.08	0.22	0.11	1.31	38.76	12.14	57.07
7	0.04	4.55	0.83	0.07	0.15	0.10	1.02	52.45	9.94	45.29
8	0.04	4.55	0.67	0.07	0.15	0.09	0.87	69.32	9.44	43.04
9	0.02	4.57	0.51	0.09	0.21	0.08	0.72	70.82	9.07	41.51
10	0.07	4.70	1.15	0.09	0.28	0.12	1.31	33.93	12.24	57.55
11	0.06	5.10	1.08	0.09	0.26	0.11	1.25	39.61	13.05	66.61
12	0.15	5.10	2.62	0.14	0.52	0.16	2.67	16.35	17.91	91.32
13	0.03	4.70	0.49	0.06	0.12	0.08	0.71	94.05	8.88	41.81
14	0.03	4.25	0.64	0.07	0.16	0.09	0.82	67.77	8.01	34.08
15	0.07	4.70	1.14	0.08	0.23	0.11	1.31	40.13	11.54	54.28

Notes: Means are per product x machine x day; "prob of stockout " "prob of recharge" is the mean probability that the brand stocks-out and gets recharged, respectively, in a day x machine; "Nb of price changes" is the number of times we observe a price change for a certain product in our dataset (during the 4 years covered by the data); "Between recharges" counts the mean number of days ("nb of days"), men number of sales ("nb of sales"), and mean revenue between product recharge event for a certain brand x machine x day.

Table 3: Pairwise diversion ratios

		To														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
From	1	0.00	0.18	0.19	0.07	0.03	0.14	0.07	0.10	0.15	0.19	0.13	0.13	0.13	0.13	0.27
	2	0.03	0.00	0.11	0.01	0.00	0.02	0.05	0.03	0.03	0.06	0.04	0.05	0.07	0.04	0.03
	3	0.01	0.03	0.00	0.00	-0.01	0.01	0.02	0.03	0.00	0.00	0.01	0.01	0.02	0.01	0.02
	4	0.08	0.04	0.06	0.00	0.27	0.36	0.14	0.14	0.08	0.08	0.04	0.04	0.02	0.22	0.12
	5	0.01	0.01	0.00	0.06	0.00	0.20	0.04	0.02	0.00	0.01	0.01	0.02	0.00	0.03	0.05
	6	0.03	0.03	0.08	0.08	0.24	0.00	0.02	0.13	0.03	0.07	0.06	0.02	0.07	0.07	0.06
	7	0.01	0.06	0.10	0.02	0.04	0.01	0.00	0.06	0.03	0.05	0.05	0.02	0.11	0.07	0.01
	8	0.01	0.02	0.09	0.02	0.02	0.08	0.06	0.00	0.02	0.04	0.03	0.02	0.05	0.09	0.05
	9	0.03	0.04	0.03	0.02	0.00	0.03	0.03	0.03	0.00	0.03	0.03	0.01	0.06	0.01	0.02
	10	0.06	0.12	0.05	0.04	0.01	0.10	0.11	0.14	0.07	0.00	0.09	0.05	0.12	0.01	0.05
	11	0.04	0.08	0.06	0.01	0.00	0.08	0.09	0.06	0.04	0.08	0.00	0.09	0.11	0.01	0.02
	12	0.09	0.22	0.22	0.05	0.08	0.05	0.11	0.13	0.04	0.12	0.24	0.00	0.14	0.10	0.15
	13	0.01	0.04	0.04	0.00	-0.01	0.03	0.07	0.04	0.03	0.03	0.03	0.01	0.00	0.01	0.00
	14	0.01	0.06	0.03	0.04	0.01	0.06	0.08	0.09	0.03	0.00	0.01	0.02	0.02	0.00	0.07
	15	0.07	0.06	0.15	0.03	0.07	0.07	0.01	0.13	0.04	0.03	0.01	0.06	-0.01	0.12	0.00

Notes: Structure-free estimated diversion ratios; lost sales from line product diverted to column product.

Table 4: Machine stockouts and machine monthly revenues

	Total machine revenue per month		
	(1)	(2)	(3)
Stockout	901.06***	1016.53***	2.35**
	(34.14)	(30.29)	(0.92)
Sales			4.74***
			(0.00)
N	8275	8275	8275
Machine FE	no	yes	yes
Month and Year FE	no	no	yes

Notes:

Table 5: Margins and the probability of a product being stocked-out in a machine x day

	Stockout of a product in a machine X day			
	(1)	(2)	(3)	(4)
Margin	1.07***	-0.31***	-0.17***	-0.29***
	(0.01)	(0.01)	(0.01)	(0.01)
Total Demand		0.06***	0.06***	0.06***
		(0.00)	(0.00)	(0.00)
weekend			-0.01***	-0.01***
			(0.00)	(0.00)
N	1823789	1823789	1823789	1823789
Manufacturer FE	no	no	yes	yes
Machine FE	no	no	no	yes

Notes: (i) Margin is the 8.5% of the price of the product; (ii) OLS regressions where the left-hand size variable is equal to 1 if the product j was out of stock in machine m and day t; (iii) Total demand is sales plus lost sales due to stock out and Weekend indicates whether the day of the week is either a Saturday or a Sunday.

Table 6: Monthly frequency of stockouts in machine and number of tobacco stores in radius around the machine

	(1)	(2)	(3)
Density of tobacco stores	-0.05*	-0.08**	-0.09**
	(0.03)	(0.03)	(0.03)
Total Demand		0.00***	0.00***
		(0.00)	(0.00)
Machine recharge frequency			-0.00**
			(0.00)
N	6748	6748	6748

Notes: (i) One observation is one x machine x month x year;
(ii) probit regressions and clustered errors at the machine level

Table 7: Product margin and duration (in days) between product's recharges

	Duration between recharges		
	(1)	(2)	(3)
Margin	-1046.89***	-727.89***	-886.79***
	(5.26)	(5.36)	(4.52)
Total Demand		-13.52***	-5.84***
		(0.05)	(0.05)
N	2904972	2904972	2904972
Machine FE	no	no	yes

Notes: (i) One observation is one product x machine x day; (ii) dependent variable is number of days between two recharges at the product level.

Table 8: Product margin and the probability a stocked out product gets re-stocked when the machine gets recharged

	Product got recharged = 1		
	(1)	(2)	(3)
Margin	4.78***	3.07***	3.20***
	(0.03)	(0.12)	(0.12)
N	528880	37377	37377
Machine FE	no	no	yes

Notes: (i) One observation is one product x machine x day conditional on at least one product in the machine getting recharged;(ii) dependent variable is dummy indicating whether the product got recharged (1) or not (0); (iii) column 1: subsample of machines that got recharged; column (2) and (3): subsample of machine sthat got recharged and products that were out of stock; (3) includes machine fixed effects.

Table 9: Demand model estimates with consumer heterogeneity

		mean	sd
Price (α_i)	Estimate	-0.260	0.324
	Standard Error	0.089	0.077
Group market share (σ_i)	Estimate	0.530	0.345
	Standard Error	0.178	0.081
N	1728045		

One observation is one product-machine-day

Table 10: Difference between counterfactual and factual Consumer and Manufacturer Surplus

	% Diff	Std. Dev.
Consumer surplus	0.190	0.250
Manufacturer surplus		
Overall mean	0.007	0.140
Manufacturer A	0.092	0.014
Manufacturer B	0.212	0.068
Manufacturer C	-0.136	0.012
Manufacturer D	-0.019	0.017