Price Promotion Holidays: Theory and Evidence

Benjamin Hansen*†
University of Oregon, NBER, IZA

Kendall Houghton*
University of Oregon

Keaton Miller*
University of Oregon

Caroline Weber*
University of Washington

June 2018

Abstract

Many retail industries feature price promotion holidays: periods close to calendar anchors when most firms simultaneously offer discounts, such as back-to-school sales in August or Black Friday sales after Thanksgiving. Models of sales by individual firms rely on supply side mechanisms, such as cost or inventory shocks, to explain discounting behavior. We analyze a price promotion holiday in the recreational marijuana industry and document facts that conflict with these traditional models. Despite an arguably exogenous increase in demand due to the proximity to a cultural event and constant marginal costs and inventories, firms choose to lower prices substantially. To explain these observations, we present a stylized model of price competition in which firms compete for customers with heterogeneous price sensitivities. An exogenous demand shock increases demand from both consumers types, and crucially increases the ratio of price-sensitive to price-insensitive consumers in a firms (firm version of choice set? bit above marginal cost?) leading to a different optimal pricing strategy for the firm.

JEL Code: D8
Keywords: Price Promotions, Nash Equilibrium, Regression Discontinuity in Time, Interrupted Time Series, Marijuana

*University of Oregon, Eugene OR, 97403-1285. Hansen: bchansen@uoregon.edu; Houghton: khoughto@uoregon.edu; Miller: keatonm@uoregon.edu; Weber: cweber5@uoregon.edu

†

1 Introduction

Demand analyses often rely upon variations in price to identify elasticities (Hausman, Leonard & Zona (1994);Berry, Levinsohn & Pakes (1995)). Standard methods assume price changes are at least partially driven by cost or inventory shocks on the supply side and that the demand system is constant across periods and prices. However, many price changes occur as part of a price promotion holiday: a period of time linked to changes in demand conditions in which retailers throughout a particular industry offer discounts on multiple products simultaneously. Well known examples of this phenomenon include back-to-school sales and Black Friday.

We present stylized facts from the recreational marijuana industry which observes a price promotion holiday on April 20th, National Weed Day (NWD), of each year. Using a comprehensive transaction-level dataset on prices, quantities, inventories, and costs for each retail firm, we show that despite constant marginal costs and minimal inventory buildup, firms choose to offer substantial discounts on NWD. Crucially, NWD existed in marijuana consumption culture before the legalization of marijuana, and firms which choose not to offer discounts still experience increased demand, suggesting that these discounts occur in the presence of an exogenous positive demand shock.

We offer a potential solution to this puzzle by presenting a model of demand for marijuana with consumers who have heterogeneous preferences over price. In the model, an exogenous shock increases demand for all types of consumers, but the ratio of price-sensitive to price-insensitive consumers changes. As a consequence, firms choose to reduce their prices. We provide suggestive evidence in favor of this mechanism by examining how the distribution of transaction sizes varies on NWD. During the promotion period, this distribution shifts towards a larger number of smaller transactions. The model also suggests that competition effectively reduces firms' ability to extract surplus from this demand shock, and indeed we document that firms earn only modest profits from NWD despite a significant increase in revenue.

Our results suggest that price promotions may result from changes in demand conditions, and thus estimates of demand elasticities which are identified through variation in prices, even in industries in which consumer stockpiling is minimal, are likely to be biased.

Beyond our presentation of the model, our stylized facts contribute to the overall understanding of price promotion holidays. These events have been difficult to study due to the scope of the promotions – as a large number of firms offer discounts across a variety of products for a short time period, capturing the behavior requires comprehensive, high-frequency sales data on products which are and are not discounted, as well as data from firms which do not offer discounts at all. Without this data, previous work on price promotion holidays has largely been limited to qualitative investigation of consumer motivation for shopping on that day (Boyd Thomas & Peters, 2011; Bell, Weathers, Hastings & Peterson, 2014; Kwon & Brinthaupt, 2015).

This data limitation has also driven previous investigators of price promotions to focus on the promotion of individual brands within a store. Promotion of one brand takes sales away from another (Papatla & Krishnamurthi (1996), Sun, Neslin & Srinivasan (2003), van Heerde, Gupta & Wittink (2003)). Promotion of one brand takes sales away from another (Papatla & Krishnamurthi (1996), Sun et al. (2003), van Heerde et al. (2003)). There is a clear short term increase in sales of the brand, but the long run effect on sales is debated (Papatla &

Krishnamurthi (1996), Mela, Gupta & Lehmann (1997), G. Dekimpe, Hanssens & Silva-Risso (1998), Pauwels, Hanssens & Siddarth (2002)). This short term increase is decomposed into cross-brand substitution, intertemporal substitution, and new demand(Gupta (1988), Bell, Chiang & Padmanabhan (1999), (Hendel & Nevo, 2006b)). The latest work suggests each component is a third of the bump in sales (van Heerde, Leeflang & Wittink, 2004).

When the focus is broadened to the overall effect for the retailer, price promotions on select brands in a store do not appear to increase traffic to the store, store-wide sales, or firm value (G. Dekimpe et al. (1998), Nijs, Dekimpe, Steenkamps & Hanssens (2001), Pauwels, Silva-Risso, Srinivasan & Hanssens (2004), Srinivasan, Pauwels, Hanssens & Dekimpe (2004)). The few studies that evaluate profit find a decrease in profit, even while revenue increases (Pauwels et al. (2002), Ailawadi, Harlam, César & Trounce (2006)). Although the timing of these brand promotions across retailers is beginning to be studied (Guyt & Gijsbrechts, 2014), no work has considered the price promotion holiday.

Hendel & Nevo (2006a) show that demand during price promotions can be linked to stockpiling behavior on the part of consumers. While we do not have panel data on purchases by households, we explore this hypothesis in our context by examining sales at the daily level both before and after NWD. We find no evidence of anticipation by consumers and a small dip in sales lasting less than a week after the promotion, which suggests that consumer stockpiling behavior is minimal.

In Section 2, we provide context on consumer holidays and NWD. In Section 3, we discuss Washington State's administrative data on recreational marijuana sales and describe retailer discounting behavior on NWD. In Section 5, we characterize discounting practices by retailer and examine whether the retailer or wholesaler pays for the discount. In Section

4, we present our main estimates on the impact of these discount decisions and placebo analyses. We present a stylized model of price and reputation competition consistent with our results in Section 6 Finally, we conclude in Section 7 and discuss the implications of our findings, external validity for other retail industries, and avenues for future work.

2 Background

Price promotion events exist for a number of industries. Some of these events exist near or on pre-established cultural and religious holidays, such as Valentine's Day, St. Patrick's Day, and Cinco De Mayo. Retailers have used advertising to add a commercial aspect to these days. Some industries, without a pre-existing holiday have sought to make them. Examples of this include Small Business Saturday, Record Store Day, National Tortilla Chip Day, National Slurpee Day, and National Donut Day. Other industries have capitalized on more obscure cultural holidays like Pi Day.

Retailers first began referring to the day after Thanksgiving as "Black Friday" in Philadelphia in the 1950's, and the first advertisement referencing it was placed in the American Philatelist in 1966 (Zorin, 2008). National use of the term spread throughout the United States during the 1980's, when retailers used it as an advertising opportunity (Pruitt, 2015). Retailers earned 7.9 billion on Black Friday 2017, and retailers offered average discounts of approximately 30% on clothes, 50% on computers and electronics, and 60% on travel (Brooke, 2017).

This paper focuses on National Weed Day, which occurs on April 20th of each year. April 20th, or 4/20, is the largest celebration of the occurrence of the number 420, an important number in marijuana culture. There are many accounts of the history of 420 in marijuana culture, with ardent purveyors of each account. The most popular story features a group of teenagers who would meet at 4:20 each day to consume marijuana and search for a legendary lost plot of marijuana. A Coast Guard member was said to have planted the patch before being deployed. These teenagers began to refer to any instance of marijuana consumption with the code 420. A member of The Grateful Dead was friends with the teenagers' brothers and the term passed along to him. Once The Grateful Dead had hold of the term, it spread throughout the marijuana consuming community (Grisso (2017), Holland (2017)). Music concerts, festivals in parks, and protests are organized throughout the world on this day (Sanger (2017), Borchardt (2017)).

Within the legalized market, NWD is associated with heavily discounted products in retail stores and increased consumption. Figure 1 plots the average price of products sold on the 20th of each month for the entire retail market. The average price on April 20th is noticeably less than the other months. The second panel of Figure 1 plots the average price of products sold on each day of April. The average price of goods sold on the 20th of the month is lower than the other days of the month.

2.1 Legal Recreational Marijuana in Washington

Possession of marijuana was made illegal in Washington State in 1923 by House Bill 3 (Camden, 2017), and illegal in the United States by the Marijuana Taxation Act of 1983 (Sanna, 2014, p. 88). In November 2012, Washington voters approved Initiative 502, which legalized the production, sale, possession, and consumption of marijuana for recreational purposes for adults over 21. The legal market opened on July 8, 2014. Although marijuana consumption was illegal in the intervening years, state law had slowly relaxed the intensity of enforcement. In 1971, possession of less than 40 grams was reduced to a misdemeanor (Dills, Goffard & Miron, 2016). Medical marijuana was decriminalized in 1998 (Gerber, 2004).

Production and retail participants in the legal market are required to document their actions in a "traceability" system that tracks every gram of marijuana produced and every dollar transferred. Random site audits of cultivators and processors, who represent the production side, and retailers are performed to ensure reporting accuracy. Misreporting firms are subject to penalties, which include fines, inventory destruction, and loss of license. We provide more details on the administrative data gathered from this system in Section 3.

3 Data and Discounts

We use an extract of the recreational marijuana administrative records maintained by the Washington State Liquor and Cannabis Board (WSLCB). These records describe each step in the marijuana supply chain. Every supply-side participant in the industry is required to report all required data to the WSLCB. Random in-person audits of retail and wholesale locations by WSLCB employees ensure data accuracy as non-compliance with data collection is punishable with inventory seizure and destruction. The final dataset follows the planting and harvesting of cannabis plants, the processing of these plants into usable goods, the sale of these goods to retailers, and the subsequent sale to consumers.

For each consumer transaction, we observe the date of purchase, and the quantity, weight, and price of items purchased. Transactions that include more than one production batch have been split into multiple entries so that each entry can include production batch specific information (we refer to these entries as transactions for the sake of clarity in the following anaylsis). A production batch is defined as a unique combination of processor, producer, harvest batch, strain, type, and size. A harvest batch is the subset of a harvest of plants that a processor has set aside to make one specific product. We define a product as a unique combination of processor, producer, strain, type, and size. In the data, a product is a time-series of product-batches. We allow the data to remain at the production-batch-transaction level to account for unobserved heterogeneity between these batches.

For the remaining analysis we refer to the processor, who takes harvested marijuana and processes it into usable form, as the brand. We refer to a retail location as a retailer. Our data suggest retail locations purchase product separately. We observe more than 82 million transactions between retailer and customers and over 2 million transactions between brands and retailers. Our data spans from the start of the legal market, July 2014, through June 2017. We observe 431 retailers, and 914 brands.

We perform minor cleaning steps on the extract. None of this cleaning significantly changes our results. Additional details, and a description of the entirety of the WSLCB administrative records are available in the Online Data Cleaning Appendix.[1]

3.1 Product Discounts

Much of our analysis relies on whether a product was discounted or not, and the extent of this discount. As such, we first determine the discount applied to every product sold on NWD. Each item in a store is associated with a production batch, and every item from

[1]Data cleaning appendix available at http://www.keatonmiller.org/s/data-cleaning-appendix.pdf

this production batch has the same characteristics (as described above). We find the modal selling price for each production batch, using transactions that occur in the sixty days before NWD.[2] This price is compared to the modal selling price for the production batch on NWD to determine the discount for this product. We allow for the discount to be zero and negative. Selection of the modal price allows idiosyncratic discounting for specific transactions to not effect our estimate of the marked product price. Broad analysis of scanner data suggests product pricing is at the modal level (Hendel & Nevo, 2006a). Figure 2 is a histogram of discounts detected at the product level.

After determining the discount applied to each product in a store, we characterize the store-level discounting practice. Our preferred store-level statistic compares the price of a weighted basket of goods on NWD to the price of the same weighted basket two weeks before NWD (Moschini, 1995). Weights are assigned to each product by the percent of total consumer transactions that product represents in the two weeks before NWD. Store level discount D is calculated using the formula $PriceIndex_{NWD} = (1 - D) \times PriceIndex_{usual}$. Negative discounts can be interpreted as a markup on NWD, which could be a strategy to respond to increased demand on this day. Before generating the price index, we remove outliers by dropping product discount observations outside the .5 and 95.5 percentiles. We perform robustness checks using the average discount of products, the percent of the store discounted 10% or more, percent of the store discounted 20% or more, the modal discount, the maximum discount, and the whether a store offers a discount above or below the average store discount in their ten-mile market. We define the ten-mile market in Section 4.

[2] We repeat this process using the sixty days before and after, and only the sixty days after to determine the selling price outside of NWD. These alternate specifications do not significantly change our results.

3.2 Product Types

Recreational marijuana product is classified into ten types: Flower, Solid Marijuana Infused Edible, Liquid Marijuana Infused Edible, Marijuana Extract for Inhalation, Marijuana Infused Topicals, Marijuana Mix Packaged, Marijuana Mix Infused, Marijuana Capsules, Marijuana Tinctures, and Marijuana Suppositories.

Flower is the least processed marijuana that is suitable for consumption. Marijuana can be infused into edible goods, both liquid and solid, which are referred to as Liquid Marijuana Infused Edible and Solid Marijuana Infused Edible, respectively. Marijuana Extract for Inhalation refers to concentrates of marijuana chemical compounds, like THC and CBD. Marijuana Infused Topicals process these chemical concentrates so that consumers can apply them topically to their skin. Marijuana Mix Packaged is flower that has been prepared for consumption without use of specialty equipment. Marijuana Mix Infused is the same as Marijuana Mix Packaged with the addition of flavor infusions (e.g. vanilla, cherry). Marijuana Capsules are chemical concentrates in pill form. Marijuana Tinctures are liquids infused with marijuana, and Marijuana Suppositories are chemical concentrates meant for vaginal or rectal insertion. Capsules, Tinctures, and Suppositories all comprise less than .1% of our sample and so we choose to omit them.

4 The Impact of Discounting

In this section, we first consider the effect of NWD on profit, revenue, and transactions for the average firm. We decompose the increase in consumption on NWD into purchases of discounted and non-discounted products, and characterize the intertemporal smoothing by consumers. We evaluate if NWD is profitable for the average firm, after controlling for this intertemporal smoothing. We then estimate the NWD effect on profit, revenue, and transactions as a function of the depth of the discount. Section 4.2 examines the robustness of our findings.

Figure 4 plots the daily average profit, revenue, and transactions over time averaged across firms. We have removed day of week and day of month effects for all. Each NWD and the surrounding days have been pooled together. Variable profit is 1.8 times the average day on NWD. The average number of transactions is 2.6 times average, and the average revenue is 2.1 times the average day. As mentioned previously, consumers are aware of the timing of NWD and are thus able to postpone purchases until NWD. Likewise, sophisticated consumers can choose to stockpile by purchasing more on NWD and consuming this excess over the following days or weeks. We do not detect a anticipatory dip, except perhaps for transactions, but there is dip after NWD that indicates stockpiling. We evaluate this more rigorously in Section 4.1. Discounted products are 19% of transactions, 18% of revenue, and 15% of NWD profit, see Figure 5.

4.1 By Discount Group

We next consider how these NWD effects on profit, revenue, and transactions vary by the level of discount offered by the store. We split retail locations into deep and shallow discount groups and see if the NWD effects are different for these groups. We define deep discounters as retail locations that discount more deeply than half of the retail locations in their ten mile market. We remove retail locations with no competitors from their ten mile market from the sample.

We estimate these effects using an augmented local linear donut regression discontinuity in time framework (Hausman & Rapson (2017); Barreca, Lindo & Waddell (2015)). The augmented local linear approach allows us to fit day of week and day of month fixed effects using the entire three year panel of data, and then conduct regression discontinuity analysis over a narrow period of time around NWD. These fixed effects control for variation in demand based on the day of the month and the day of the week. We estimate these fixed effects using

$$Log(Y)_t = \beta_0 + \sum_{j=1}^{6} \delta_j (DOW_j)_t + \sum_{j=1}^{30} \gamma_j (DOM_j)_t + \sum_{j=1}^{N} \alpha_j (Holiday_j)_t + e_t \tag{1}$$

where DOW_j represents each day of the week, except for Tuesday, which is left as our baseline. DOM_j represent the day of the month. We choose to omit the 5th of the month as the baseline. $Holiday_j$ represents each holiday in the three year panel (e.g. Christmas, Fourth of July). We include $Holiday_j$ to allow our fixed effects to not absorb variation in demand caused by idiosyncratic holiday placement in time. After estimating β_j and α_j, we remove the effects from our panel of data. The subsequent regression discontinuity analysis is performed on the residual. For more information on our use of augmented local linear approach, please see the Online Appendix.

Our window of analysis for the donut regression discontinuity in time spans the two months before and two months after NWD. We conduct robustness analysis of this choice of window width in Section 4.2.

After removing seasonal fixed effects, our analyses use the following template:

$$
\begin{aligned}
Log(Y)_{rt} =& \beta_0 + \beta_1 Pre_t \times T_t + \sum_{j=1}^{3} \sigma_j (NWD-j)_t + \beta_2 NWD_t \\
&+ \sum_{j=1}^{7} \alpha_j (NWD+j)_t + \beta_3 Post_t + \beta_4 Post_t \times T_t \\
+DD_r \Big[& \delta_0 + \delta_1 Pre_t \times T_t + \sum_{j=1}^{3} \lambda_j (NWD-j)_t + \delta_2 NWD_t \\
&+ \sum_{j=1}^{7} \phi_j (NWD+j)_t + \delta_3 Post_t + \delta_4 Post_t \times T_t \Big] + e_{rt}
\end{aligned}
\tag{2}
$$

where our dependent variable Y is revenue, transactions, or profit for day t at retailer r. $Pre_t \times T_t$ is a trend line for before NWD. We include indicators for the three days before NWD, $(NWD-j)_t$ where $j \in 1,2,3$. We then add an indicator for NWD, NWD_t, and a set of seven indicators for the week after NWD, $(NWD+j)_t$ where $j \in 1,2,...,7$. We allow for an intercept shift after NWD, $Post$, and a new trend, $Post_t \times T_t$. Each of these terms is allowed to be different for the deep discount group, where DD_r is an indicator variable for deep discounters. We allow the deep discount group to have an initial intercept difference as well by including DD_r as a variable on its own. A deep discounter is a retailer that offers a larger store-level discount than half of retailers in its ten mile market. The ten mile market is the set of all retailers within time miles of the retailer in question.

Figures 6 through 8 plot the raw data and fit line for our three dependent variables. Circles are raw observations with seasonal trends removed, where the observations have been averaged by discount group and day. The solid lines are the fit lines from our estimation. The light gray line and hollow circles correspond to deep discounters. The black line and filled in circles correspond to shallow discounters. The estimates that correspond to these figures are reported in Table 2 as well. *Anticipation* is the sum of coefficients attached to the indicator variables for the three days before NWD. *Recovery* is the sum of the coefficients attached to the indicator variables for the week after NWD.

Profit increases by \$2,595 on NWD. Retailers have an additional 466 transactions on this day and earn \$8,459 more revenue. The deep discounting group have 194 more transactions on average, an approximately 40% increase of the non deep discounters' 466 transactions, and make \$2,672 more in revenue, an approximately 30% increase of the non deep discounters' \$8,459 of revenue. Notably, deep discounting firms do not earn more profit than non deep discounting firms on NWD, on average. Although more customers are buying from the retailer and the retailer is earning more revenue, these do not translate to increased profit. This could be explained by the lower profit margins for retailers that are offering deeper discounts.

We do not detect an anticipatory drop in profit, transactions, or revenue in the days before NWD, but we do observe a dip in all of these after. Retailers earn \$1,772 less profit than normal in the week following NWD, have 182 fewer transactions, and earn \$5,067 less revenue. These do not differ for the deep discounting group. We calculate the net profit, transactions, and revenue using the spike on NWD and the dip in the following week. The average firm earns \$824 from NWD, has 284 more transactions, and earns \$3,392 more revenue. These transaction and revenue numbers are higher for deep discounting firms, following the increases listed above.

We explicitly test for parallel trends using the coefficient on $Pre-Trend$. We do not detect a statistically different trend for the deep discounting group compared to the non deep discounting, although there is a different pre-treatment intercept for transactions.

NWD improves firm profitability on average, but deeper discounters are not rewarded with higher profits. Our analysis of the post-NWD trends indicates that deep discounters experience a statistically significantly different post-NWD profit trend, but this trend is barely economically significant - the deep discounters earn $2 more per day on average.

4.2 Robustness

Tables 2 and 3 and Figures 9 and 10 provide evidence that our preferred estimate is robust to a variety of specifications. We discuss robustness for the profit specification, for evaluation of revenue and transactions, see the Online Appendix.

Table 2 repeats our profit specification from Table 1 for different measures of store discount level. 10 $Mile\ Mean$ classifies deep discounters as those who discount more deeply than the average store-level discount in their ten-mile market. Max classifies a retail location as a deep discounter if its maximum discount is deeper than half of the maximum retailer discounts in the 10-mile market. $Mode$ finds the modal discount offered by a retail store and classifies a deep discounter as one whose modal discount is larger than half of their 10-mile market. 10% and 20% measure the percent of the store 10 or 20 percent off, respectively, and classify deep discounters as those that have a larger percent of the store meeting this threshold than half of their ten-mile market. The first column is our specification from Table 1. We see that the NWD estimate is consistent. All other coefficients are similar as well.

Figure 9 varies the portion of the market that is classified as a deep discounter. In each of our previous specifications, a retailer is classified as deep discounter if it offers a discount larger than half of its ten mile market, that is if it offers a discount in the top half. Here we vary how close to the deepest discounter in their ten-mile market a retailer must be in order to be classified as a deep discounter. The x-axis records if a retail location must be in the top half, third, fourth, fifth, or sixth of discounters. The y-axis plots the point estimate for one of six key variables of interest. We include bands for the 95% confidence interval. Clockwise, beginning on the top left, we plot the coefficient on NWD, post-NWD intercept shift, and the post NWD trend change for the baseline. We plot the trend shift, intercept shift, and NWD coefficients for deep discounters. All coefficient estimates appear robust to variation in the deep discount threshold.

Figure 10 reconsiders the geographic area that a retailer's market is calculated over. Our main estimation classifies deep discounters as those with discounts in the top half of their ten mile market. Figure 10 shows the point estimates for our key variables as we vary the size of the market. The x-axis records the size of the market, 5-mile, 10-mile, 15-mile, 20-mile, and 25-mile. The y-axis records the value of the point estimate. We include bands for the 95% confidence interval. Clockwise, beginning on the top left, we plot the coefficient on NWD, post-NWD intercept shift, and the post NWD trend change for the baseline. We plot the trend shift, intercept shift, and NWD coefficients for deep discounters. All coefficient estimates appear robust to variation in the size of the geographic market.

Figure 3 evaluates the robustness of our preferred specification to different augmented local linear approaches. We vary the period of time over which are fixed effects are fit. The first column is our preferred specification from Table 1. The last column is a non-augmented local linear approach, in which we fit the fixed effects over the four month period of our primary regression. All estimates are consistent across the different fixed effect time periods.

5 Who Discounts and What Does It Cost Them

We next consider why retailers are offering the combination of discounts they are (hereafter referred to store level discounts), where we define a retailer as a retail location. We observe considerable variation in the store level discounts offered on NWD (see Figure 3), as measured by the price index. Additionally, we observe multiple NWDs for the same retailer. In order to understand the source of this variation, we estimate the store level discount as a function of observables. Given the potential for collinearity between many of these measures, we first estimate the relationship between each observable and our outcome variable separately using

$$Log(Discount)_{rt} = \beta_0 + \beta_1 X_{rt} + e_{rt} \tag{3}$$

where an observation is the store level discount, as measured by the price index, for retailer r on NWD in year t. X_{rt} is one of six observables - $Rank$, $Log(Profit)$, $Comp$, $Log(Time)$, Log(Transactions), $andLog(Weight)$. $Rank_{rt}$ is the measure of store rank in the ten mile market, where rank is determined by total revenue, and the "ten mile market" is all other retail stores within ten miles of the retailer in question. The largest seller in a market has $Rank_{rt} = 1$, the second largest has $Rank_{rt} = 2$, and so on. $Log(Profit)_{rt}$ is log of the average variable profit for the store in the ten weeks before NWD. $Number\ of\ Comp_{rt}$ is the number of a competitors in the ten mile market. $Log(Time)_{rt}$ is the log of the number of weeks a retail location has been open. $Log(Transactions)_{rt}$ is the average weekly number of transactions in a store, and $Log(Weight)_{rt}$ is the weekly average weight of usable marijuana sold. We allow the standard errors to cluster at the retailer level.

In Column (7), we include all measures together and in Column (8) include a retailer fixed effect. After including a retailer fixed effect, the remaining identifying variation comes from changes in observables and discount behavior for a given retailer over time.

We present the results in Table 4. $Log(Transactions)_{rt}$ and $Log(Weight)_{rt}$ are meaningfully related to the overall discount level in the store. Although when both measures are included, the coefficient attached to $Log(Weight)_{rt}$ is no longer significant. This suggests that much of the power of $Log(Weight)_{rt}$ is coming from its ability to proxy for the number of transactions a store usually experiences in a week. A one percent increase in the average number of transactions in a week leads to a seven percent increase in discount depth.

In Figure 11, we plot the average daily wholesale price for products that are sold to consumers at a discount on NWD and those that are sold at the normal price. Wholesale prices are not lower for discounted products.

6 A Stylized Model of Discounts

In this section, we develop a stylized model of supply and demand with an aim of replicating the stylized facts documented throughout the previous sections in a parsimonious way. In the model, two firms compete for consumers by setting prices. The key tradeoff is intertemporal: firms have a reputation which evolves according to the share of consumers captured each period. National Weed Day offers firms an opportunity to change their reputation more significantly than on other days. Therefore, a discount offered today increases the firm's reputation and increases the firm's ability to capture profits tomorrow.

Formally, consider two firms denoted by a and b, who compete in a market with two periods with time denoted by $t = 1, 2$. These firms offer a single product each, with constant marginal costs c_a, c_b. Firms enter the first period with a state variable, which we call reputation, determined exogenously and denoted by x_a and x_b. In each period, firms choose their own price p_a or p_b.

Consumers are identical and choose a single firm (or neither) in each period. The choice specific utility for consumer i and firm j is given by

$$u_{ij} = \beta x_j - \alpha p_j + \epsilon_{ij}.$$

Let ϵ_{ij} be i.i.d. as Type I Extreme Value. Then the share of consumers choosing firm j is

$$s_j = \frac{\exp(\beta x_j - \alpha p_j)}{1 + \sum_k \exp(\beta x_k + \alpha p_k)}.$$

The firm's state x'_j evolves from period 1 to period 2 as a function of the share of consumers they are able to capture in the first period s_j:

$$x'_j = \gamma x_j + (1 - \gamma) s_j.$$

We characterize the equilibrium in this game by backwards induction. In the second period, the x_j are fixed and so firms are simply playing a Nash-Bertrand game. The profit for firm j in the second period is given by

$$\pi_j^2(p_j; x'_j, x'_{-j}, p_{-j}) = (p_j - c_j) s_j.$$

By taking the first order condition with respect to prices and rearranging, we can obtain an expression that characterizes the optimal price in period 2:

$$
\begin{aligned}
0 &= s_j + (p_j - c_j)\frac{\partial s_j}{\partial p_j} \\
&= s_j - \alpha(p_j - c_j)s_j(1 - s_j) \\
&= s_j(1 + \alpha(p_j - c_j)(1 - s_j)) \\
p_j &= \frac{1}{\alpha(1 - s_j)} + c_j
\end{aligned}
$$

To characterize the solution in the first period, first write $\pi_j^*(x'_j, x'_{-j}) = \max \pi_j^2(p_j; x'_j, x'_{-j}, p_{-j})$. Then the problem facing firm j in period 1 can be written as

$$
\max_{p_j} (p_j - c_j)s_j(p_j; p_{-j}, x_j, x_{-j}) + \sigma\pi_j^*(x'_j(\cdot), x'_{-j}(\cdot)
$$

While a subgame perfect equilibrium exists, it cannot be found analytically. Instead, we solve for the equilibrium numerically. We solve for the equilibrium as a function of the initial reputations x_a, x_b for reasonable values of $\alpha, \beta, c_a, c_b, \lambda$, and σ. The results are shown in Figure 12.

7 Conclusion

Price promotion holidays are a feature of many industries. These holidays, periods of time in which retailers offer discounts simultaneously, have been difficult to study given the need for extensive retail and wholesale data for multiple firms. We use a novel dataset from the recreational marijuana industry for Washington State, for which these data are available and in which the industry celebrates a price promotion holiday on April 20th of year.

We empirically examine four questions: What discounts do firms offer on NWD? Why do they offer these discounts? Do retailers or brands pay for these discounts? How do consumers respond? Using stylized facts from this analysis, we construct a two-stage model of firm discounting behavior.

Store level discounts vary widely. The average store discount is 5% and the standard deviation is 11%. Some retailers choose to mark up products on NWD, as a way to capture profit from the increased demand. Observables do not explain why stores choose to offer a certain overall discount level, although increased average weekly transactions are associated with deeper discounting, on average. Product discounts cost retailers more than firms.

Consumer appear to stockpile products, but do not appear to wait to purchase until NWD. Although first-time, non-repeat consumers may choose to select NWD as a day to consumer the product. Consumers purchase both discounted and non-discounted products on NWD. We find no evidence that profit varies by discount depth, although transactions and revenue do. Even after accounting for the dip in profit in the week after, NWD is profit improving for the average firm. Although not largely so, the average firm earns $824, which is roughly a quarter of an average day's profit. Our results are robust to a variety of specifications.

These results contribute to the understanding of strategic price promotion behavior. The lack of change in post-NWD profit and the lack of variation in profit on NWD by discount group is suggestive of retailers who are best responding in equilibrium. The net profit is positive for the average firm, which suggests the existence of NWD is welfare enhancing for firms. In contrast, a day of increased marijuana consumption may be welfare decreasing for consumers if there are negative externalities to marijuana consumption.

References

Ailawadi, K. L., Harlam, B. A., César, J., & Trounce, D. (2006). Promotion Profitability for a Retailer: The Role of Promotion, Brand, Category, and Store Characteristics. *Journal of Marketing Research*, *43*(4), 518–535.

Barreca, A. I., Lindo, J. M., & Waddell, G. R. (2015). Heaping-Induced Bias in Regression-Discontinuity Designs. *Economic Inquiry*.

Bell, D. R., Chiang, J., & Padmanabhan, V. (1999). The Decomposition of Promotional Response: An Empirical Generalization. *Marketing Science*, *18*(4), 504–526.

Bell, G. C., Weathers, M. R., Hastings, S. O., & Peterson, E. B. (2014). Investigating the Celebration of Black Friday as a Communication Ritual. *Journal of Creative Communications*, *9*(3), 235–251.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, *63*(4), 841–890.

Borchardt, D. (2017). 17 of the craziest 420 cannabis celebration events. *Forbes*.

Boyd Thomas, J. & Peters, C. (2011). An exploratory investigation of Black Friday consumption rituals. *International Journal of Retail & Distribution Management*, *39*(7), 522–537.

Brooke, Z. (2017). Infographic: Black friday 2017 by the numbers. *American Marketing Association Marketing News Weekly*.

Camden, J. (2017). Marijuana making its way around circle from legal to illegal to legal. *The Spokesman-Review*.

Dills, A. K., Goffard, S., & Miron, J. (2016). Dose of Reality: The Effect of State Marijuana Legalizations. *Policy Analysis*, (799).

G. Dekimpe, M., Hanssens, D. M., & Silva-Risso, J. M. (1998). Long-run effects of price promotions in scanner markets. *Journal of Econometrics*, *89*(1-2), 269–291.

Gerber, R. J. (2004). *Legalizing Marijuana: Drug Policy Reform and Prohibition Politics*. Greenwood Publishing Group.

Grisso, J. (2017). The real story behind why "4/20" became a stoner holiday. *Mother Jones*.

Gupta, S. (1988). Impact of Sales Promotions on When, What, and How Much to Buy. *Journal of Marketing Research*, *25*(4), 342–355.

Guyt, J. Y. & Gijsbrechts, E. (2014). Take Turns or March in Sync? The Impact of the National Brand Promotion Calendar on Manufacturer and Retailer Performance. *Journal of Marketing Research*, *51*(6), 753–772.

Hausman, C. & Rapson, D. S. (2017). Regression discontinuity in time: Considerations for empirical applications. Technical report, National Bureau of Economic Research.

Hausman, J., Leonard, G., & Zona, J. D. (1994). Competitive Analysis with Differentiated Products. *Econometrics of Imperfect Competition*, (34), 159–180.

Hendel, I. & Nevo, A. (2006a). Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica*, *74*(6), 1637–1673.

Hendel, I. & Nevo, A. (2006b). Sales and consumer inventory. *The RAND Journal of Economics*, *37*(3), 543–561.

Holland, B. (2017). The hazy history of ”420”. *History*.

Kwon, H. J. & Brinthaupt, T. M. (2015). The motives, characteristics and experiences of US Black Friday shoppers. *Journal of Global Fashion Marketing*, *6*(4), 292–302.

Mela, C. F., Gupta, S., & Lehmann, D. R. (1997). The Long-Term Impact of Promotion and Advertising on Consumer Brand Choice. *Journal of Marketing Research*, *34*(2), 248.

Moschini, G. (1995). Units of measurement and the stone index in demand system estimation. *Agricultural & Applued Economics Association*, *77*(1), 63–68.

Nijs, V. R., Dekimpe, M. G., Steenkamps, J.-B. E., & Hanssens, D. M. (2001). The Category-Demand Effects of Price Promotions. *Marketing Science*, *20*(1), 1–22.

Papatla, P. & Krishnamurthi, L. (1996). Measuring the dynamie effects of promotions on brand choice. *Journal of Marketing Research*.

Pauwels, K., Hanssens, D. M., & Siddarth, S. (2002). The Long-Term Effects of Price Promotions on Category Incidence, Brand Choice, and Purchase Quantity. *Journal of Marketing Research*, *39*(4), 421–439.

Pauwels, K., Silva-Risso, J., Srinivasan, S., & Hanssens, D. M. (2004). New Products, Sales Promotions, and Firm Value: The Case of the Automobile Industry. *Journal of Marketing*, *68*(4), 142–156.

Pruitt, S. (2015). What's the real hisotry of black friday. *History*.

Sanger, B. (2017). 10 best 420 celebrations around the world. *Herb*.

Sanna, E. (2014). *Marijuana: Mind-Altering Weed*. Simon and Schuster.

Srinivasan, S., Pauwels, K., Hanssens, D. M., & Dekimpe, M. G. (2004). Do Promotions Benefit Manufacturers, Retailers, or Both? *Management Science*, *50*(5), 617–629.

Sun, B., Neslin, S. a., & Srinivasan, K. (2003). Measuring the Impact of Promotions on Brand Switching When Consumers Are Forward Looking. *Journal of Marketing Research*, *40*(4), 389–405.

van Heerde, H. J., Gupta, S., & Wittink, D. R. (2003). Is 75% of the Sales Promotion Bump Due to Brand Switching? No, Only 33% Is. *Journal of Marketing Research*, *40*(4), 481–491.

van Heerde, H. J., Leeflang, P. S. H., & Wittink, D. R. (2004). Decomposing the Sales Promotion Bump with Store Data. *Marketing Science*, *23*(3), 317–334.

Zorin, E. (2008). 'black friday'–why and when? *Chicago Tibune*.

8 Figures and Tables

Figure 1: Comparison of April 20th Prices to 20th of Other Months

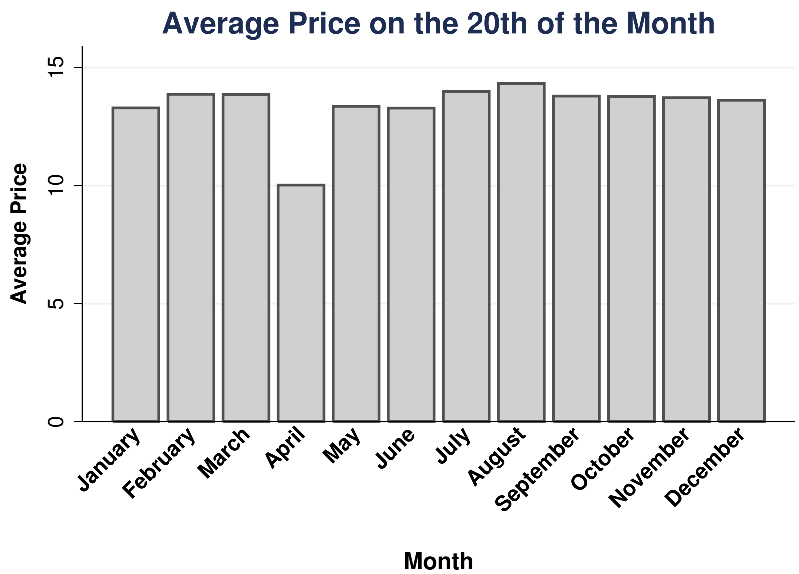

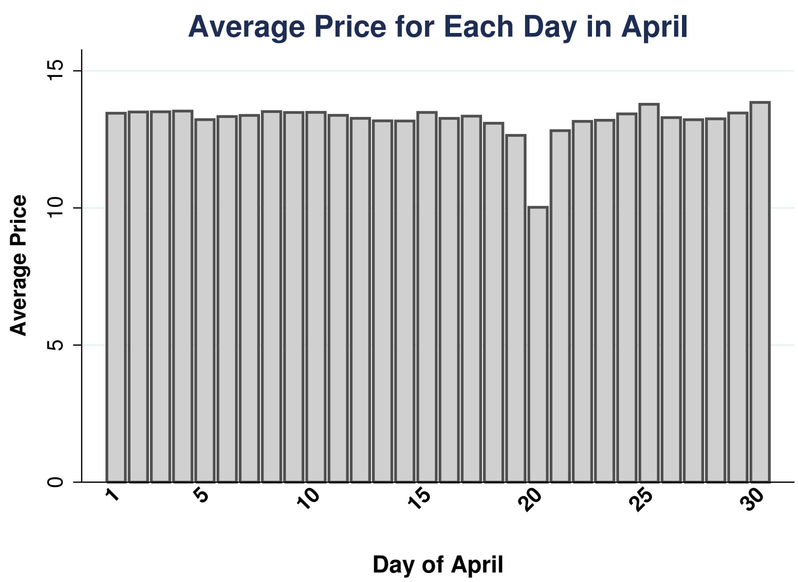

This plot compares the prices in stores on April 20th to the prices in stores on the 20th of the other months. We find the average price of products sold in a store on the 20th of the month, and then average these store prices across the industry. The "other months" bar pools every 20th of the month in our time series together, with the exception of April. The "April" bar is a pool of the three observed April 20ths. For products sold by the gram, we use the price per gram. For other products, we use the transaction price for the item. The store price is an average taken over transactions, and therefore frequently bought products are weighted more heavily and non-purchased products are dropped from the calculation.

Figure 2: Product Discounts

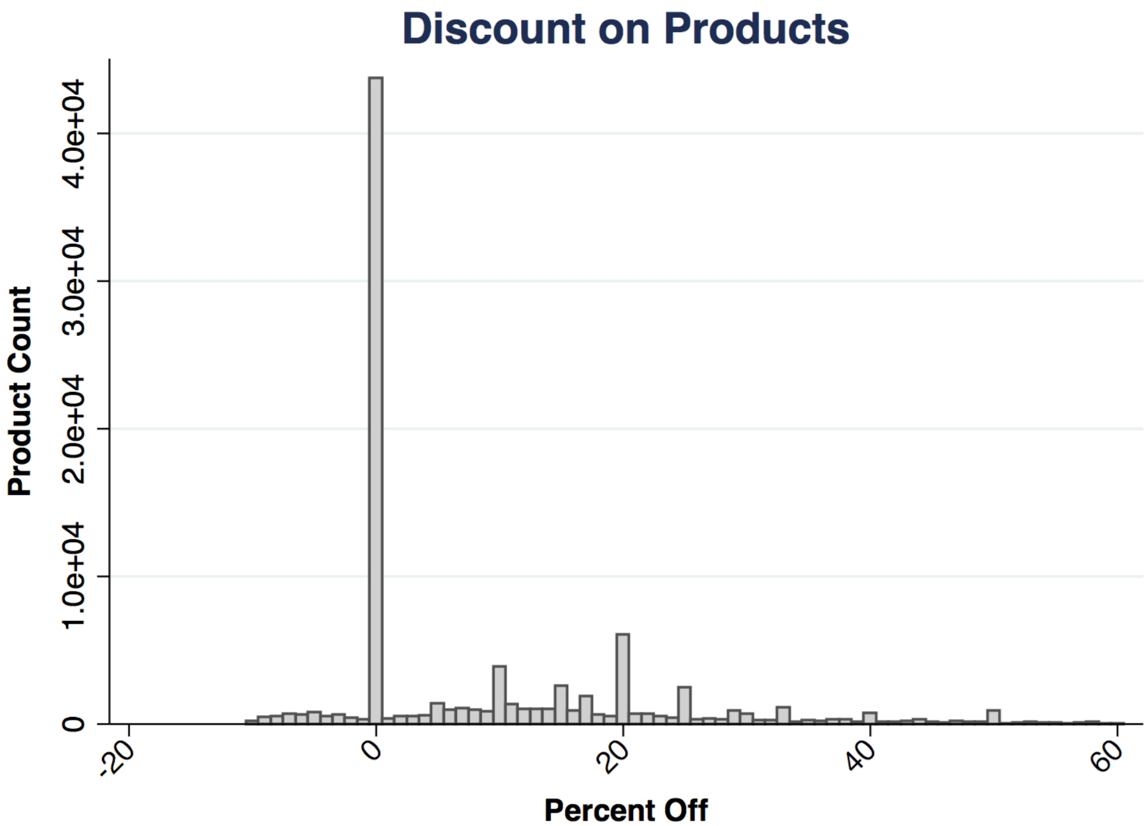

We plot the frequency of discounts at the product level, for all products sold on NWD and sold at least once before this date. To calculate the discount, we find the modal price on transactions of the product for the preceding two weeks and compare this price to modal price of transactions for the product on NWD. The modal price is used in order to account for idiosyncratic discounting that would affect the average price. Negative discounts indicate a markup is being placed on the product on NWD. We limit our observations to those within the .5 and 95.5 percentiles of the distribution. Products that were not purchased on NWD are dropped from our sample, as our products that are sold for the first time on NWD.

Figure 3: Retailer Discount

We aggregate individual product discounts to create a measure of the overall store discount on NWD. The above plot shows the frequency of different store level discounts. We generate a price index using the frequency of purchases of products for the preceding two weeks, we then recalculate the cost of the same bundle of products on NWD. Store level discount D is calculated using the formula $PriceIndex_{NWD} = (1 - D) \times PriceIndex_{usual}$. We then multiply D by 100 in order to translate it into percent terms. Negative discounts can be interpreted as a markup on NWD. Before generating the price index, we remove outliers by dropping product discount observations outside the .5 and 95.5 percentiles.

Figure 4: Retailer Variable Profit, Transactions, and Revenue Over Time, Post-NWD Dip

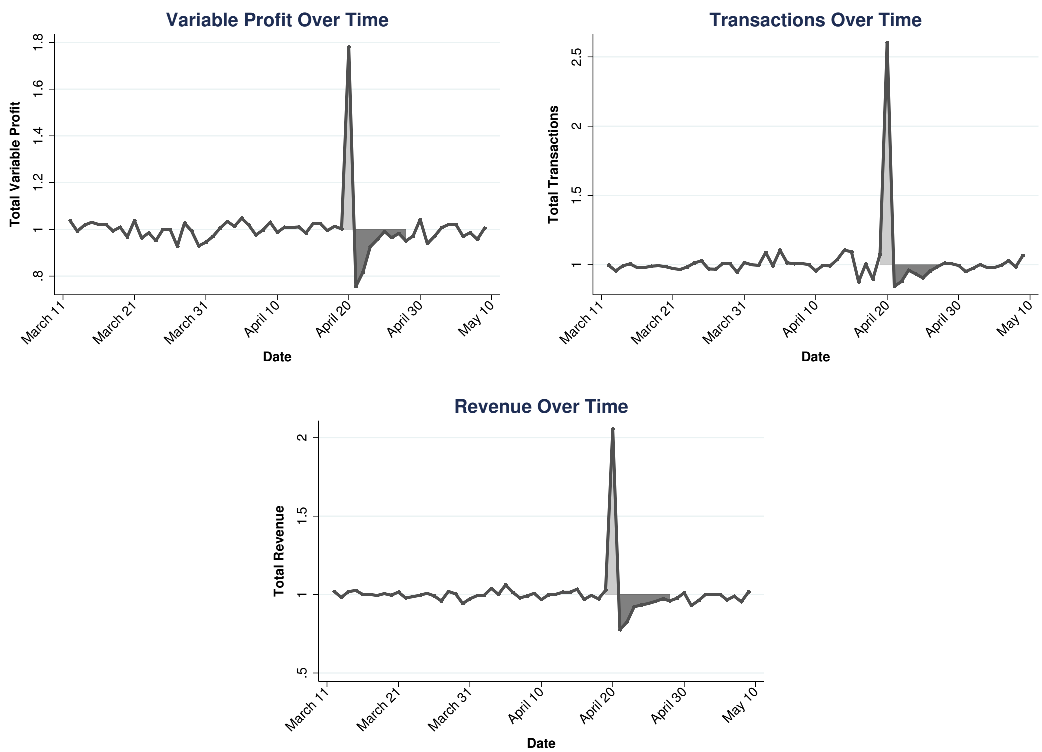

The above plot shows variable profit over time, in thousands, with time marked in days before or after NWD, as well as transactions and revenue. We have removed day of month and day of week effects, and generated the average profit, transactions, and revenue for a retailer in the industry on each day. We have pooled all years. For product sold in grams, profit is the tax exclusive price per gram less the price per gram paid in wholesale cost, multiplied by the weight of the product. For other products, we use the tax exclusive revenue from the sale of the item minus the wholesale cost of the item. When plotting, we shift the line to match the mean of all observations before removing day of week and day of month effects. The light gray shading indicates the difference in spike on NWD and the pre-NWD mean. The dark gray shading indicates the difference between the observed average on the seven days after NWD and the pre-NWD mean.

Figure 5: Retailer Variable Profit Over Time, Discounted vs. Non-Discounted

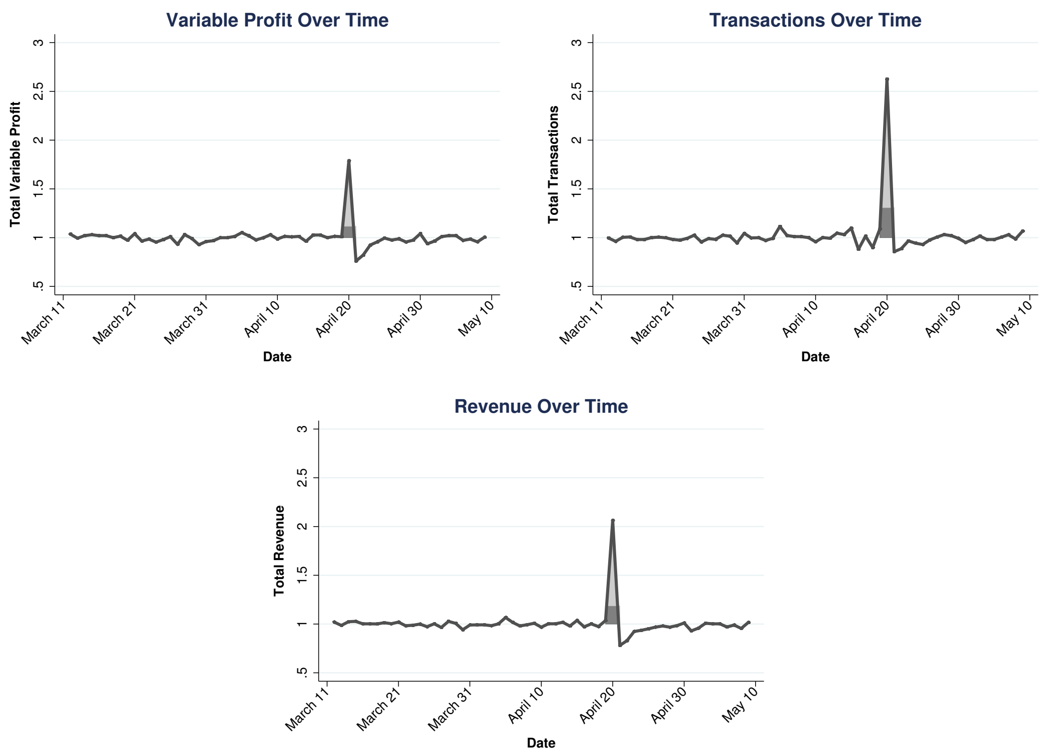

The above plot shows variable profit over time, in thousands, with time marked in days before or after NWD, as well as transactions and revenue. We have removed day of month and day of week effects, and generated the average profit, transactions, and revenue for a retailer in the industry on each day. We have pooled all years. For product sold in grams, profit is the tax exclusive price per gram less the price per gram paid in wholesale cost, multiplied by the weight of the product. For other products, we use the tax exclusive revenue from the sale of the item minus the wholesale cost of the item. When plotting, we shift the line to match the mean of all observations before removing day of week and day of month effects. The light gray shading indicates the difference in spike on NWD and the pre-NWD mean. The dark gray shading indicates the portion of this spike that is due to discounted products.

Figure 6: Changes in Variable Profit

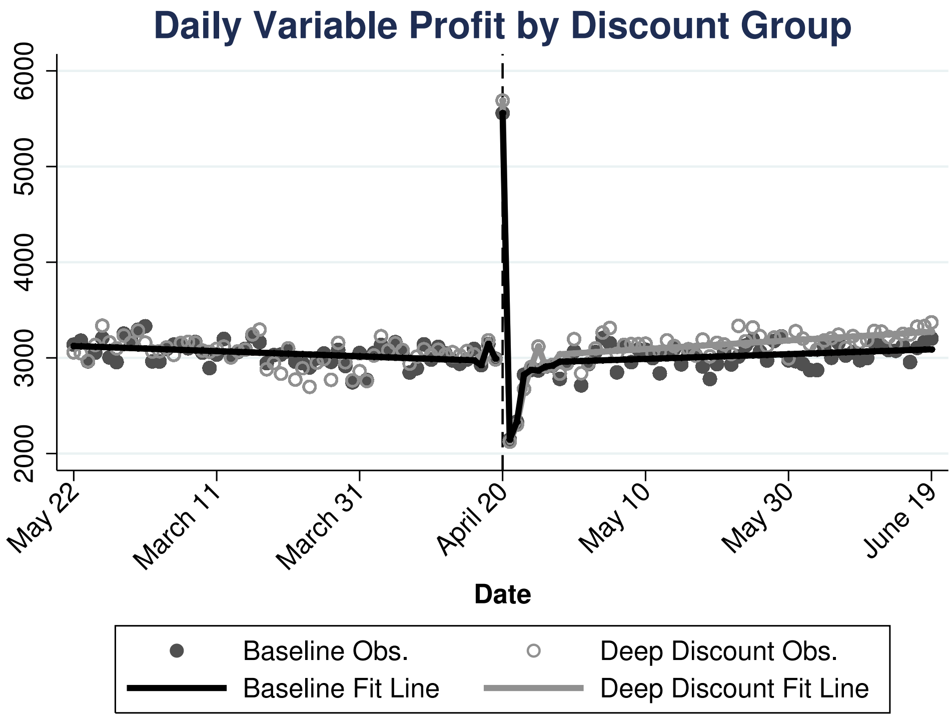

The plot above shows the fit line and observed data for retailer daily variable profit. We have separated this estimation into two groups - shallow and deep discounters. Dots represent the observed data, averaged across retailers in a discount group for each day. Shallow discounters are represented by the black line and grey dots. Deep discounters are represented by the black line with dark gray dots. We have removed day of week and day of month fixed effects when representing the observed data and fit line. We shift the line upward to match the mean of all observations without day of week and day of month effects removed. The February dates in 2016 are one day ahead of the marked time due to leap year.

Figure 7: Changes in Revenue

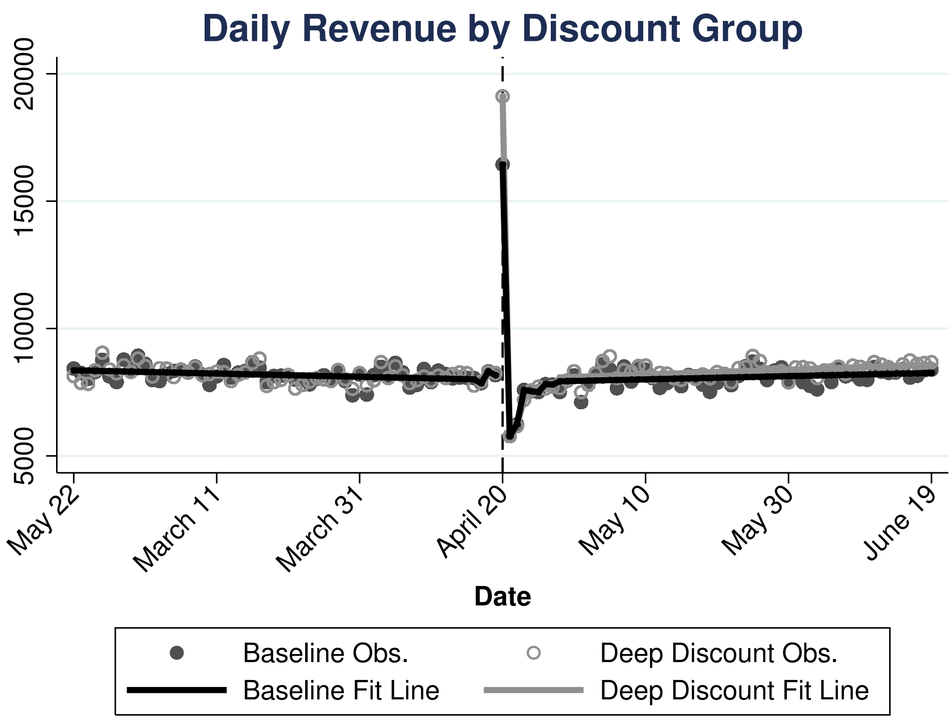

The plot above shows the fit line and observed data for retailer daily revenue. We have separated this estimation into two groups - shallow and deep discounters. Dots represent the observed data, averaged across retailers in a discount group for each day. Shallow discounters are represented by the black line and grey dots. Deep discounters are represented by the black line with dark gray dots. We have removed day of week and day of month fixed effects when representing the observed data and fit line. We shift the line upward to match the mean of all observations without day of week and day of month effects removed. The February dates in 2016 are one day ahead of the marked time due to leap year.

Figure 8: Changes in Transactions

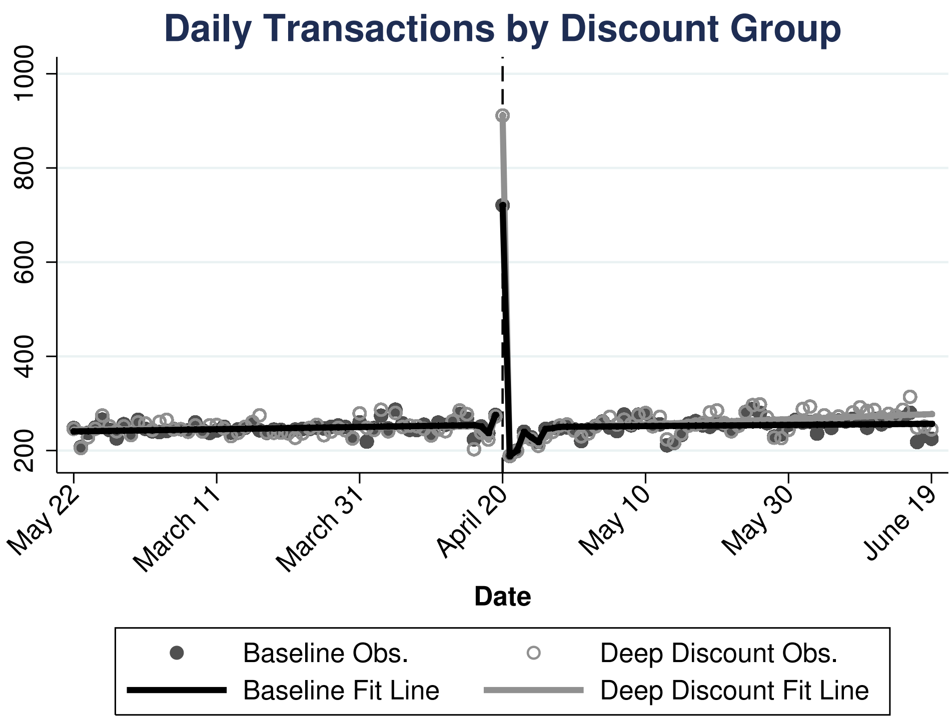

The plot above shows the fit line and observed data for retailer daily transactions. We have separated this estimation into two groups - shallow and deep discounters. Dots represent the observed data, averaged across retailers in a discount group for each day. Shallow discounters are represented by the black line and grey dots. Deep discounters are represented by the black line with dark gray dots. We have removed day of week and day of month fixed effects when representing the observed data and fit line. We shift the line upward to match the mean of all observations without day of week and day of month effects removed. The February dates in 2016 are one day ahead of the marked time due to leap year.

Figure 9: Deep Discounter Robustness

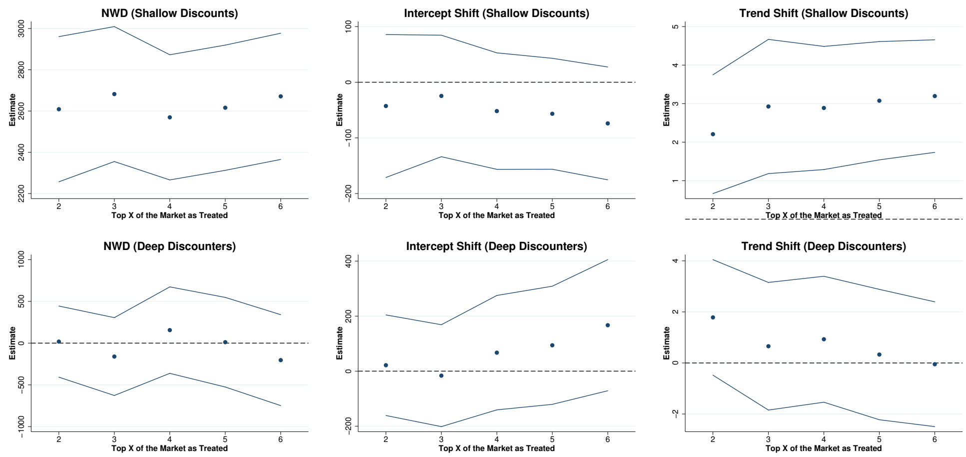

In each of our previous specifications, a retailer is classified as deep discounter if it offers a discount larger than half of its ten mile market, that is if it offers a discount in the top half. Here we vary how close to the deepest discounter in their ten-mile market a retailer must be in order to be classified as a deep discounter. The x-axis records if a retail location must be in the top half, third, fourth, fifth, or sixth of discounters. The y-axis plots the point estimate for one of four key variables of interest. We include bands for the 95% confidence interval. Clockwise, beginning on the top left, we plot the coefficient on post-NWD intercept shift for the baseline specification, and the post NWD trend change for the baseline. We plot the trend shift for the deep discounters, and the intercept shift for deep discounters.

Figure 10: Market Geographic Distance Robustness

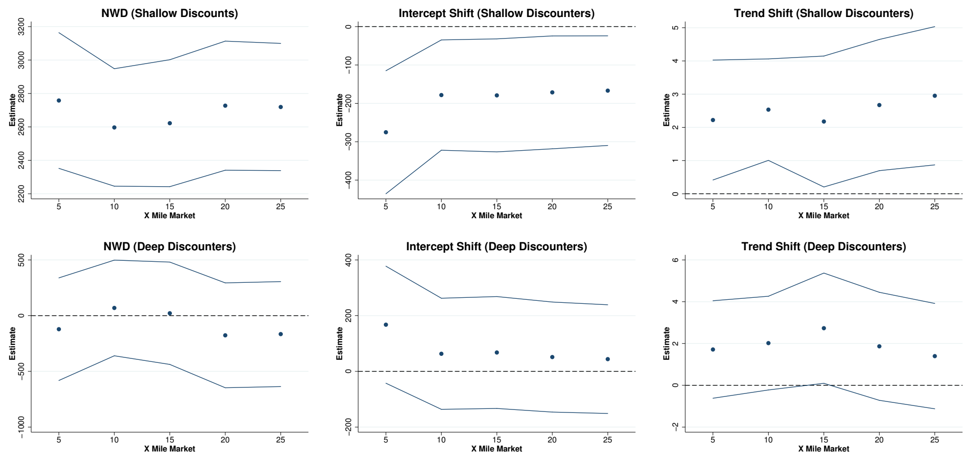

Our main estimations classify deep discounters as those with discounts in the top half of their ten mile market. The above plot shows the point estimates for our key variables as we vary the size of the market we consider. The x-axis records the size of the market, 5-mile, 10-mile, 15-mile, 20-mile, and 25-mile. The y-axis records the value of the point estimate. We include bands for the 95% confidence interval. Clockwise, beginning on the top left, we plot the coefficient on post-NWD intercept shift for the baseline specification, and the post NWD trend change for the baseline. We plot the trend shift for the deep discounters, and the intercept shift for deep discounters.

Figure 11: Wholesale Price of Discounted and Non-Discounted Products

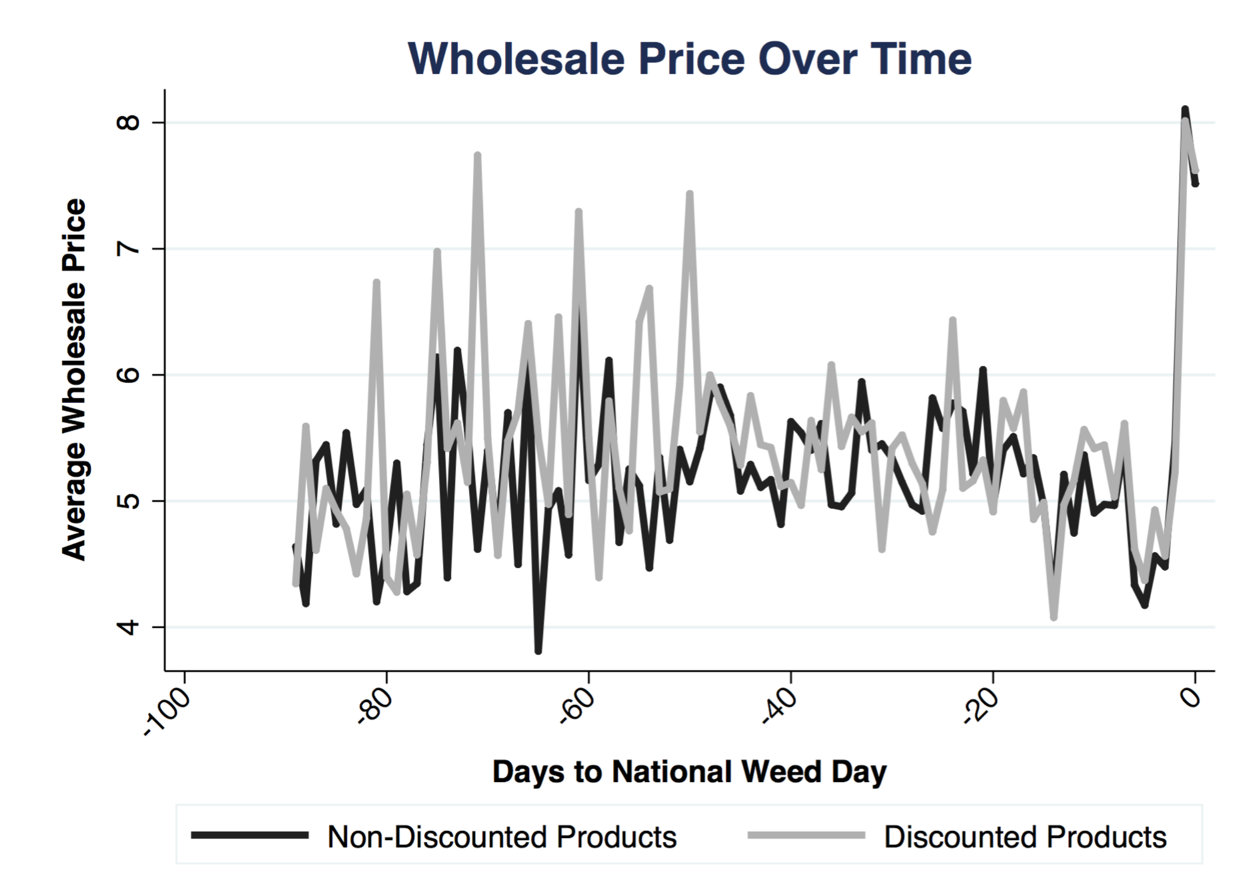

We plot the daily average wholesale price for discounted and non-discounted products. Discounted products are those sold on NWD to consumers at a discount. Non-discounted are those sold to consumers on NWD at the normal price. We do not observe a meaningful difference in wholesale price between the two groups. Prices spike for both in the days immediately preceding NWD.

Figure 12: Model outputs as a function of initial firm reputation

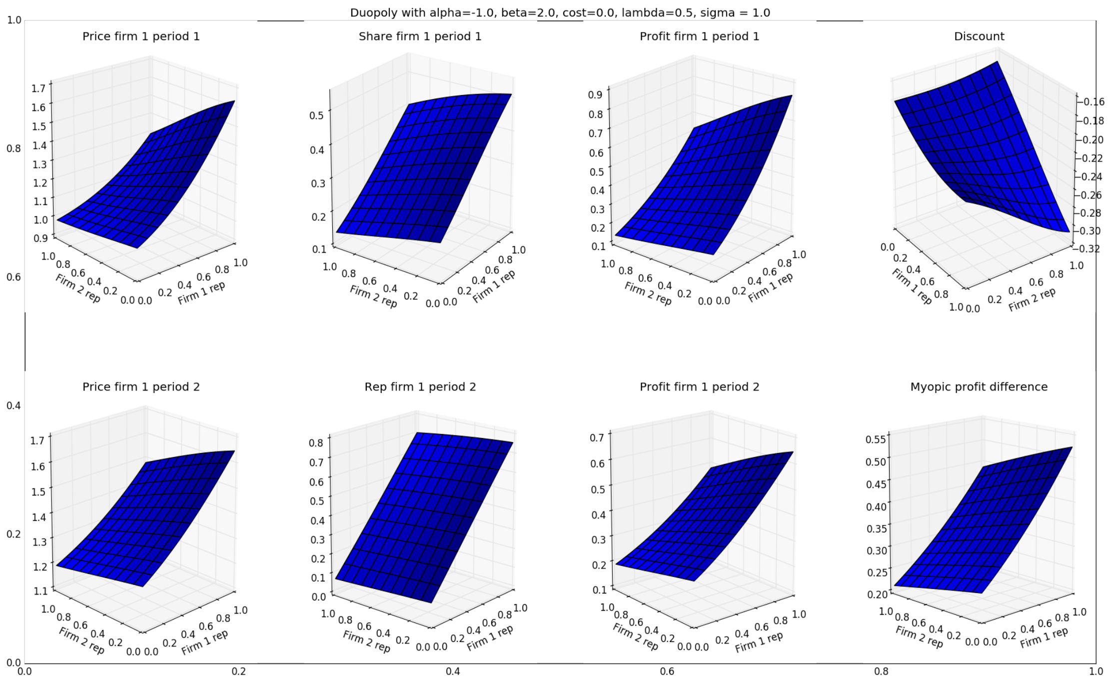

These plots illustrate equilibrium outputs of the model described in Section 6.

Table 1: Changes in Variable Profit, Transactions, and Revenue

	Variable Profit *(dollars)*	Transactions *(unit)*	Revenue *(dollars)*
Baseline (Shallow Discount)			
Pre-Trend	-2.684**	0.238**	-6.281**
	(-2.38)	(2.06)	(-2.10)
Anticipation	155.2	0.621	352.8
	101.5	12.73	289.9
NWD	2595.4***	465.9***	8458.8***
	(14.06)	(9.83)	(13.61)
Recovery	-1772.2***	-181.8***	-5067.1***
	210.1	23.88	563.5
Intercept Shift	-23.57	-5.377	-103.1
	(-0.55)	(-1.07)	(-0.95)
Post-Trend	2.489***	0.122	6.303***
	(3.64)	(1.54)	(3.44)
Deep Discount ×			
Intercept	197.9	73.91***	832.1
	(0.85)	(2.81)	(1.41)
Pre-Trend	-0.0704	-0.110	0.131
	(-0.03)	(-0.60)	(0.03)
Anticipation	95.08	-22.27	115.2
	217.5	18.77	450.8
NWD	134.9	193.6***	2672.9***
	(0.59)	(2.75)	(2.85)
Recovery	-406.9	-29.18	-1423.9
	356.6	38.90	916.6
Intercept Shift	65.93	-1.530	121.6
	(0.69)	(-0.20)	(0.79)
Post-Trend	2.030*	0.417**	3.306
	(1.74)	(2.51)	(1.01)
Constant	-592.5***	-120.2***	-1975.3***
	(-2.66)	(-4.25)	(-3.22)
Observations	67729	68454	67743

The above table describes our main estimation results. We fit day of month and day of week fixed effects using the augmented local linear method. *Anticipation* is the sum of the coefficients attached to to the three indicator variables for the three days before NWD. *Recovery* is the sum of the coefficients attached to the seven indicator variables for the week after NWD.

Table 2: Deep and Shallow Discounter Robustness

	(1) Price Index	(2) 10-Mile Mean	(3) Max	(4) Mode	(5) 10%	(6) 20%
Baseline(Shallow Discount)						
Pre-Trend	-2.684**	-1.975*	-2.060	-3.742**	-2.052*	-2.365*
	(1.209)	(1.192)	(1.396)	(1.716)	(1.065)	(1.280)
Anticipation	155.2.144	-64.248	21.368	25.470	-24.392	46.171
	(102.461)	(115.591)	(120.146)	(169.934)	(117.359)	(124.094)
NWD	2595.393***	2773.439***	2511.895***	2862.111***	2835.304***	2864.147***
	(14.057)	(193.580)	(182.683)	(189.884)	(195.155)	(184.971)
Recovery	-1772.214	-1857.699	-1535.054	-2145.474	-1836.882	-1841.860
	(210.140)	(217.025)	(321.026)	(289.000)	(328.336)	(333.648)
After NWD Intercept Shift	-23.572**	-177.735**	-131.306	-146.466**	-92.728	-97.075
	(-0.552)	(75.288)	(81.010)	(70.363)	(65.502)	(72.018)
Post-Trend	2.489***	3.462***	3.577***	3.167***	3.354***	3.610***
	(3.639)	(0.870)	(1.112)	(0.752)	(1.154)	(1.188)
Deep Discount $\times$						
Intercept	197.891	-80.233	555.673**	-340.869	-336.561	-254.945
	(0.854)	(228.331)	(264.577)	(214.170)	(294.227)	(232.919)
Pre-Trend	-0.0704	-1.568	-1.210	1.947	-1.333	-0.803
	(-0.03)	(2.069)	(1.964)	(2.220)	(1.807)	(1.808)
Anticipation	95.08	140.036	-14.572	-27.015	70.812	-55.267
	(217.483)	(213.383)	(242.731)	(225.878)	(249.985)	(252.044)
NWD	134.904	-269.008	223.195	-446.091*	-378.889	-430.672*
	(0.59)	(221.408)	(252.699)	(226.870)	(276.571)	(242.577)
Recovery	-406.915	-357.241	-981.426	203.022	-384.407	-377.195
	(356.548)	(336.975)	(435.716)	(366.801)	(429.819)	(434.111)
Intercept Shift	65.931	60.307	-7.589	6.908	-97.159	-93.395
	(0.694)	(100.317)	(108.917)	(116.633)	(99.227)	(99.808)
Post-Trend	2.030*	0.313	0.100	0.899	0.479	0.000
	(1.742)	(1.097)	(1.362)	(1.119)	(1.429)	(1.416)
Constant	-592.502	-252.264	-607.450**	-118.017	-122.621	-159.489
	(275.878)	(282.423)	(284.456)	(269.042)	(278.178)	(284.162)
Observations	67743	67743	67743	67743	67743	67743

Our main estimation uses a price index for products at each retail location to generate a measure of discount depth for the overall store. We repeat our preferred specification from Table 1, using alternate ways of classifying deep and shallow discounters. 10-*Mile Mean* classifies a retail location as a deep discounter if its store-level discount is greater than their 10 mile market mean. *Max* classifies a retail location as a deep discounter if its maximum discount is deeper than half of the maximum retailer discounts in the 10-mile market. *Mode* finds the modal discount offered by a retail store and classifies a deep discounter as one whose modal discount is larger than half of their 10-mile market. 10% and 20% measure the percent of the store 10 or 20 percent off, respectively, and classify deep discounters as those that have a larger percent of the store meeting this threshold than half of their 10-mile market.

Table 3: Fixed Effects Robustness

	Augmented Local Linear			Four Month Window		
	Variable Profit *(dollars)*	Transactions *(unit)*	Revenue *(dollars)*	Variable Profit *(dollars)*	Transactions *(unit)*	Revenue *(dollars)*
Baseline (Shallow Discount)						
Pre-Trend	-2.684**	0.238**	-6.281**	-3.852***	0.189	-8.691**
	(-2.38)	(2.06)	(-2.10)	(-2.77)	(1.45)	(-2.53)
Anticipation	155.2	0.621	352.8	49.20	-98.72	-545.1
	101.5	12.73	289.9	356.8	27.15	907.6
NWD	2595.4***	465.9***	8458.8***	2702.7***	466.5***	8630.9***
	(14.06)	(9.83)	(13.61)	(12.64)	(9.63)	(12.88)
Recovery	-1772.2***	-181.8***	-5067.1***	-880.3	-153.1	-3231.5
	210.1	23.88	563.5	321.7	28.29	714.7
Intercept Shift	-23.57	-5.377	-103.1	-51.10	-4.001	-171.4
	(-0.55)	(-1.07)	(-0.95)	(-1.13)	(-0.70)	(-1.61)
Post-Trend	2.489***	0.122	6.303***	3.621***	0.156	8.630***
	(3.64)	(1.54)	(3.44)	(3.33)	(1.42)	(3.32)
Deep Discount ×						
Intercept	197.9	73.91***	832.1	196.3	73.58***	826.0
	(0.85)	(2.81)	(1.41)	(0.84)	(2.86)	(1.46)
Pre-Trend	-0.0704	-0.110	0.131	-0.125	-0.116	-0.0270
	(-0.03)	(-0.60)	(0.03)	(-0.06)	(-0.65)	(-0.01)
Anticipation	95.08	-22.27	115.2	106.1	-19.22	158.9
	217.5	18.77	450.8	240.2	19.92	496.9
NWD	134.9	193.6***	2672.9***	136.9	193.9***	2680.0***
	(0.59)	(2.75)	(2.85)	(0.62)	(2.82)	(3.07)
Recovery	-406.9	-29.18	-1423.9	-393.7	-29.77	-1403.4
	356.6	38.90	916.6	356.2	40.77	862.8
Intercept Shift	65.93	-1.530	121.6	66.24	-1.151	125.6
	(0.69)	(-0.20)	(0.79)	(0.69)	(-0.15)	(0.89)
Post-Trend	2.030*	0.417**	3.306	2.043*	0.420**	3.368
	(1.74)	(2.51)	(1.01)	(1.76)	(2.39)	(0.97)
Constant	-592.5***	-120.2***	-1975.3***	2521.2***	219.6***	6790.9***
	(-2.66)	(-4.25)	(-3.22)	(9.51)	(7.23)	(9.58)
Observations	67729	68454	67743	67729	68454	67743

The three left columns are the results as reported in Table 1. We then estimate these coefficients again, but fit the day of month and day of week fixed effects within the four month window of our regression discontinuity instead of using the augmented local linear approach. These resulting coefficients are reported in the three columns on the right.

Table 4: Retailer Discount Depth by Observables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Rank	-0.037						-0.057	0.249
	(0.050)						(0.101)	(0.399)
Log(Profit)		6.585					2.097	34.688
		(4.616)					(9.913)	(24.048)
Number of Comp.			0.014				0.010	0.119
			(0.038)				(0.071)	(0.213)
Log(Time)				0.332			-0.326	-0.614
				(0.394)			(0.508)	(1.551)
Log(Transactions)					6.329***		7.476***	9.359*
					(1.083)		(1.038)	(5.171)
Log(Weight)						5.694*	-2.485	-26.064
						(3.319)	(8.056)	(21.858)
Constant	10.241***	-4.689	9.714***	8.832***	-1.972	-1.721	-2.303	-25.335
	(0.599)	(10.393)	(0.618)	(1.535)	(2.050)	(6.902)	(11.167)	(47.329)
Retailer FE?	No	No	No	No	No	No	No	Yes
Observations	640	636	660	686	684	685	634	634

Using the price index measure of retailer overall discount depth, we estimate the discount depth as a function of observables. $Rank_{rt}$ is the measure of store rank in the ten mile market, where rank is determined by total revenue, and the "ten mile market" is all other retail stores within ten miles of the retailer in question. The largest seller in a market has $Rank_{rt} = 1$, the second largest has $Rank_{rt} = 2$, and so on. $Log(Profit)_{rt}$ is log of the average variable profit for the store in the ten weeks before NWD. $NumberofComp_{rt}$ is the number of a competitors in the ten mile market. $Log(Time)_{rt}$ is the log of the number of weeks a retail location has been open. $Log(Transactions)_{rt}$ is the average weekly number of transactions in a store, and $Log(Weight)_{rt}$ is the weekly average weight of usable marijuana sold. We allow the standard errors to cluster at the retailer level. In Column (7), we include all measures together and in Column (8) include retailer fixed effect. After including retailer fixed effect, the remaining identifying variation comes from changes in observables and discount behavior for a given retailer over time.