THE UNIVERSITY OF CHICAGO

SELF-CONTROL AND OVERCONSUMPTION: EMPIRICAL EVIDENCE FROM CIGARETTE PURCHASES

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
QUOC DANG HUNG HO

CHICAGO, ILLINOIS
AUGUST 2026

Abstract

Using transaction data from more than 10,000 smokers across 121 convenience stores, this study examines how consumers exercise self-control when purchasing vice goods. I exploit variation in exposure to multipack discounts (e.g., "buy 2, get $1 off") and show that many smokers forgo these deals in favor of frequent single-pack trips. This pattern is specific to cigarettes: the same consumers redeem multi-unit discounts of other products at typical rates. Using an instrumental variable strategy, I find that buying an extra pack due to a discount causes smokers to increase daily smoking by 35%. I then estimate a mixed-type structural model that allows smokers to differ in their overconsumption risk and awareness of that risk. I find that most smokers are prone to overconsume, and about two-thirds are at least partially aware of this risk. Counterfactuals show that reintroducing a half-sized pack reduces consumption by about 2% when priced at 80% of the standard pack and by up to 4.5% when priced at 60%. By contrast, the model predicts that banning quantity discounts would reduce consumption by less than 1%, consistent with evidence from New York's statewide discount ban.

Table of Contents

List of Tables

List of Figures

Acknowledgments

I am deeply grateful to my partner, Minh Nghiem, for her unwavering support and for being my family during the hardest part of graduate school.

I also thank all the doctoral students in the Marketing program at Chicago Booth whom I have had the privilege to know over the past five years. It has been a great privilege to learn from them and to share with them the many ups and downs of being a doctoral student.

This dissertation has also benefited tremendously from the guidance of the faculty in the Marketing group at Chicago Booth, especially my committee members.

I thank Andreas Kraft for being an excellent role model for how to navigate the first few years as a junior faculty member.

I thank Jean-Pierre Dubé, whose rigorous feedback has fundamentally shaped how I think about research, and whose standards I will strive to follow throughout my research career.

I thank Pradeep Chintagunta for admitting me to the program and for his invaluable advice on how to navigate the research process.

Finally, I thank my advisor, Avner Strulov-Shlain, for being with me from the very beginning. He has been extraordinarily generous with his time, advice, and support. I cannot put into words how much his mentorship has meant to this dissertation and to my doctoral experience. This would not have been possible without him, and I am grateful to him for everything.

1 Introduction

Consumers are known to ration purchases of vice goods like cigarettes, deliberately forgoing quantity discounts in favor of smaller pack sizes (Wertenbroch 1998, Marti and Sindelar 2015). Many smokers do so based on the belief that "one smokes less with smaller pack sizes" (Gomez and Guevara 1993). These patterns suggest a form of self-control, where consumers limit how much they buy to prevent overconsumption. Yet, we still lack causal evidence that larger quantities actually lead to higher consumption, and that consumers are aware of this effect. Understanding the extent of this awareness is crucial for how vice goods are bought, sold, and regulated. If most consumers anticipate their overconsumption risk and ration their purchases, regulation may be unnecessary. Firms can profit by offering smaller, premium-priced packages. But if many consumers lack such awareness, firms have incentives to promote bulk purchasing of harmful goods through quantity discounts, indicating a need for regulatory oversight. Although self-control through purchase rationing is well-studied in theory and laboratory settings (see, for example, Mishra and Mishra 2011, Dobson and Gerstner 2010, Jain 2012, Jain and Li 2018), empirical evidence in the field remains limited.

This paper has two goals. First, it provides real-world evidence of self-control in cigarette purchases by showing that (i) having more cigarettes on hand causally increases smoking, and (ii) many smokers are aware of this risk and engage in purchase rationing. Second, it quantifies heterogeneity in self-control across consumers and explores its implications for marketing strategy and intervention design.

It is useful to distinguish between two facets of self-control in this context. The first occurs at the purchase stage, where consumers buy in smaller sizes to avoid consuming more than intended. Throughout this paper, I use the term *self-control* or *sophistication* interchangeably to refer specifically to this purchase-stage behavior. The second facet arises at the consumption stage, where consumers attempt to moderate their use of a product already purchased. I refer to *overconsumption* as a failure of self-control at this stage.

I use a novel dataset covering all transactions at 121 convenience stores in Maryland between August 2020 and December 2024. The dataset contains information on store-level prices and discounts, as well as time series of purchases made by 10,607 smokers. My empirical analysis exploits the widespread use of multi-unit discounts for cigarettes and beverages at these stores (e.g., "buy 2 get $1 off"). Variations in exposure to these deals allow me to identify behavioral patterns consistent with overconsumption and self-control.

Two data patterns are consistent with self-control through purchase rationing. First, in 89% of trips, smokers purchase only one or two packs. Notably, a substantial share of smokers almost always buy just one pack at a time. This is especially surprising because these consumers tend to restock very quickly:

over half of single-pack trips are followed by another purchase within two days. Many of these single-pack trips occur even when a multipack discount is available. Thus, even setting aside shopping costs, these consumers could have paid less by purchasing a two-pack bundle rather than making two separate single-pack purchases.[1]

Second, this tendency to forgo multi-unit discounts appears specific to cigarettes. Smokers who regularly ignore cigarette multipack discounts still take advantage of multi-unit discounts on beverages and chocolate at rates similar to other consumers. This helps rule out general discount aversion or price insensitivity as explanations for their avoidance of cigarette multipacks.

A self-control interpretation rests on the premise that smokers are vulnerable to overconsumption: they smoke more when they have more cigarettes on hand. Prior studies have documented that larger purchases are associated with greater consumption for tempting goods such as soda or cigarettes (Gordon and Sun 2015, Kim and Ishihara 2025). However, it remains unclear how much of this link is causal, and how much of it is due to unobserved demand, e.g. smokers may buy more cigarettes because they face sudden cravings and have already planned to smoke more.

To disentangle the causal effect of having more cigarettes on the smoking rate from unobserved cravings, I employ an instrumental variable strategy that leverages variations in multipack deals across store visits. Specifically, I use whether a smoker visits a store offering multipack discounts as an instrument for the number of packs purchased. The instrument is relevant as it increases the likelihood of buying two packs. Critically, it should not affect the smoking rate except through that purchase decision (exclusion restriction). Two facts support the exogeneity of this instrument. First, multipack discounts are typically store-level promotions negotiated between cigarette manufacturers and retailers. They last for extended periods and are not designed to respond to smokers' short-term cravings. Thus, over the 4-year panel, the identifying variation comes from the same smokers being observed under different store-level discount spells, both within and across the stores they visit. Second, smokers are no more likely to visit a discounted store on high-craving days, implying that discount exposure moves the number of packs purchased but does not directly affect smoking through other channels.

I find that having more cigarettes on hand leads to increased smoking. On average, purchasing an extra pack due to multipack discounts causes smokers to return in only 48% extra time, corresponding to an 35% increase in daily consumption. This result provides direct evidence of overconsumption and suggests that smokers may have an incentive to ration their cigarette purchases.

Given the severe public health consequences of smoking,[2] I turn to the question of identifying effective

[1] While consumers may buy cigarettes from other outlets not captured in the data, this does not explain why they return to purchase cigarettes again after just 1-2 days at the stores in the sample.

[2] Cigarette smoking kills half a million of U.S. adults each year and imposes billions of dollars in economic costs (CDC)

interventions for consumers at risk of overconsumption. The right intervention depends crucially on the degree of consumer sophistication. Sophisticated smokers already avoid buying in large quantities and may benefit from access to smaller pack sizes, provided those packs are priced appropriately to deter youth uptake. In contrast, if smokers underestimate their overconsumption risk, policies should reduce the appeal of purchasing multiple units, such as banning quantity discounts or taxing larger packs more heavily. Public education campaigns may also help by raising awareness of these risks.

I develop a mixed-type discrete choice model to estimate the degree of consumer sophistication. In the model, smokers choose between purchasing one or two packs of cigarettes on each shopping trip. This decision depends on two key factors: (i) current consumption need, which I assume is stochastically determined by recent consumption, and (ii) awareness of overconsumption. I incorporate overconsumption based on the framework of Bernheim and Rangel (2004). Specifically, purchasing two packs entails a probability that the smoker will consume more than their actual need and return to the store sooner. I derive this probability using the estimates from the instrumental variable analysis. Consuming more than one's need is costly, which reflects the underlying assumption that smokers inherently want to limit their cigarette intake. A key feature of the model is that consumers are heterogeneous both in whether they are prone to overconsume and in how they perceive that risk. Casual consumers do not overconsume when they have more cigarettes on hand. Among consumers who are prone to overconsume, sophisticated consumers fully anticipate this risk and therefore are more likely to avoid buying two packs. Naive consumers believe they always consume exactly what they need and therefore exhibit a stronger preference for purchasing two packs. The remaining group is partially naive: they recognize the risk of overconsumption but not to its full extent.

I find that 86% of smokers are prone to overconsume, while 14% are casual smokers. The remaining types are 37% sophisticated, 29% partially naive, and 20% fully naive. Smokers perceive the cost of smoking an unintended pack to be \$32.6,[3] while the effort cost of an extra shopping trip is about \$9.0. Thus, smokers may still have strong incentives to buy two packs even without a discount.

I then use the structural model estimates to evaluate several counterfactual policies. First, I simulate a ban on quantity discounts, which would make the price of two packs equals exactly twice the price of a single pack. The effect of this policy on consumption is minimal, with the largest reduction of only 1.9%, observed among partially naive smokers. The overall reduction in consumption is less than 1%. Next, I validate the model using a real-world policy change: New York's statewide discount ban, implemented on July 1, 2020. I evaluate this policy using a difference-in-differences analysis and find that its observed impact on consumption is also modest, consistent with the model's prediction. This alignment provides

[3]For reference, the average price of a pack of cigarettes is about \$10.8.

external validation that the structural model captures the key behavioral response to a discount ban.

Next, I consider the reintroduction of a smaller pack of 10 sticks. To mitigate concerns about youth smoking initiation,[4] I simulate a range of prices in which the half-sized pack costs between 60% and 100% of the standard 20-stick pack.[5] Although this pack is more expensive per cigarette, it gives smokers an opportunity to consume less than their planned need. Naive and casual smokers rarely choose the half pack, while sophisticated smokers take it up at much higher rates: from about 61% of purchases when it is priced at 60% of the standard pack to 19% when it is priced the same as the standard pack. This translates into overall consumption reductions ranging from about 0.8% to 4.5%, depending on the half-pack price.

The first contribution of this paper is to provide evidence of self-control in the market. I extend a distinguished literature documenting consumer demand for *commitment devices*[6] across domains like savings (Thaler and Benartzi 2004, Ashraf et al. 2006), exercise (Milkman et al. 2014, Royer et al. 2015), and addictive goods (Giné et al. 2010, Schilbach 2019, Allcott et al. 2022). However, most existing evidence comes from controlled experiments in which commitment devices are usually designed by researchers and offered to participants.[7] In contrast, my setting focuses on natural commitment devices that arise organically in the marketplace. In my context, the commitment devices take the form of forgone discounts and extra shopping trips, costs that consumers willingly incur without being nudged. Because this behavior emerges from consumer responses to real-world product and pricing structures rather than experimental framing, the findings are directly relevant to marketing practice. They highlight how firms can, whether to the benefit of consumers or not, influence consumer behavior through packaging, pricing, and promotion.

The closest related field study is Wertenbroch (1998), which finds that demand for vice products increases less in response to price reductions than demand for healthier alternatives.[8] My paper complements this in two key ways. First, my analysis focuses on quantity choices for a single product, thereby avoiding confounds from cross-product differences unrelated to health. Second, rather than relying on aggregate store-level data, which may reflect heterogeneity across consumers, I leverage individual-level data to provide more direct evidence of within-person purchase rationing.

The second contribution of this paper is to estimate the degree of consumer sophistication and examine its implications for marketing strategy and intervention design. Methodologically, I build on a rich literature

[4]Pack sizes smaller than 20 have been banned in the U.S. (Family Smoking Prevention and Control Act, 2009) and in many other countries, primarily due to concerns that smaller, more affordable packs may encourage youth uptake.

[5]Following Marti and Sindelar (2015), who find that some smokers are willing to pay substantial price premia for smaller package sizes.

[6]Commitment devices are less desirable options intentionally chosen to regulate future consumption.

[7]Important exceptions include Della Vigna and Malmendier (2006) and Zhang et al. (2022), which examine commitment through contract choices in subscription settings. In the former, consumers choose flat-rate gym memberships even when a pay-per-visit plan would be cheaper—an attempt to commit to exercising more. In the latter, users of a digital content platform opt to pay per chapter rather than subscribe monthly, presumably to curb overconsumption.

[8]The comparison was between "regular" and "light" versions of different products.

that estimates psychological parameters in choice models using field data (e.g., Conlin et al. 2007, DellaVigna et al. 2012, Hastings and Shapiro 2013, Handel 2013, Dubé et al. 2017, Joo 2023). To my knowledge, this is the first paper to estimate the Bernheim and Rangel (2004) model using consumer purchase data. The closest related paper is Miller et al. (2025), which estimates heterogeneity in consumer awareness of inertia in newspaper subscriptions. My findings highlight how firms can offer smaller pack sizes to help sophisticated consumers regulate their consumption of vice goods, and how some policies may fall short if they fail to account for self-control. In that sense, I add to the rapidly growing literature on behavioral industrial organization (e.g., Lacetera et al. 2012, Grubb and Osborne 2015, Gui and Drerup 2025, Strulov-Shlain 2023, Miller et al. 2025, Natan 2024, Kraft and Rao 2024, Huang 2024, Bao et al. 2025, Lin and Strulov-Shlain 2025).

This paper also relates to the emerging literature at the intersection of marketing and health policy, particularly studies evaluating tobacco control policies (Wang et al. 2016, Tuchman 2019, Bonfrer et al. 2020, Wang et al. 2021, Goli et al. 2022, Goli et al. 2024, Loghmani and Goli 2024, Levine 2025). My paper is one of the few to consider pack size as a potential regulatory lever.

Finally, this paper relates to the literature examining the link between purchase quantity and consumption. Prior work show that discounts can generate sales spikes without increasing consumption due to stockpiling (e.g., Erdem et al. 2003, Hendel and Nevo 2006, Ching and Osborne 2020, Alwan et al. 2024). Most of this literature focus on products with largely fixed consumption, like detergent or ketchup. A notable exception is Gordon and Sun (2015), which documents cigarette stockpiling. The key distinction is that their setting involves much lower prices (average $1.65/pack vs. $10.78/pack in mine), making bulk purchases economically feasible. Similar to other work in the endogenous consumption literature (Ailawadi and Neslin 1998, Sun 2005, Kim and Ishihara 2025), their identification for the quantity-consumption effect relies on a structural model with functional form assumptions. To my knowledge, my paper is the first to establish a causal link between purchase quantity and consumption using an instrument that exogenously shifts quantity choice.

2 Theoretical Framework

2.1 Intuition

I present a simple model to demonstrate the role of overconsumption and self-control in my setting. My framework is based on the addiction model of Bernheim and Rangel (2004). In particular, I incorporate two key properties: (1) smokers' choices and preferences may diverge, and (2) smokers may adopt a precommitment strategy.

Smokers make a binary decision: purchase a small pack of cigarettes or a large pack at a lower per-unit price. The large pack is economically appealing but it also poses the risk of overconsumption. A smoker may intend to take advantage of the multipack discount and save the extra cigarette pack for later. Yet, once having two packs on hand, they may end up consuming both packs just as quickly as they would have consumed a single one. Thus, their actual consumption path is different from their perceived consumption path at the time of purchase. Some smokers recognize this risk. Anticipating their own tendency to overconsume when more cigarettes are available, they choose the smaller, more expensive pack as a form of precommitment.

2.2 Setup

Suppose there are two time periods: $t = 1$ and $t = 2$. At the start of period 1, a smoker draws a consumption need c, which can be either low (c_L) or high (c_H). Based on this need, they decide how many cigarette packs to purchase. If $c = c_H$, the smoker needs to consume two packs in period 1, and therefore must buy both packs at $t = 1$, then return at $t = 2$ to restock. If $c = c_L$, the smoker has two options: (i) buy one pack and consume it at $t = 1$, then return at $t = 2$ to buy another; or (ii) buy two packs at $t = 1$, consume one at $t = 1$ and save the second for period 2. Without loss of generality, I normalize the utility of consuming one pack per period to zero.

I now focus on the case where the smoker draws a low consumption need ($c = c_L$). The net present cost of buying one pack in each period is simply

$$u_1 = \alpha p_1(1 + \delta) + \delta h,$$

where α is the marginal utility of money, δ is the per-period discount factor, and h is the shopping cost. To allow for overconsumption, I assume that if a smoker with $c = c_L$ buys two packs at $t = 1$, they may feel tempted to consume both packs immediately, effectively behaving as if $c = c_H$. This excess consumption due to temptation yields a long-term disutility u_0, which can be thought of as the net lifetime cost of smoking one additional pack. I assume that u_0 is incurred in period 2 to reflect temporal separation between smoking and its consequences (see, for example, Chaloupka et al. 2019).

I assume that excess consumption occurs with probability π, which is conceptually akin to the probability of entering the *hot mode* in Bernheim and Rangel (2004). In their model, a potentially addicted individual operates in two modes: (i) a “cold” mode in which they rationally choose their most preferred option; and (ii) a “hot” mode in which they make impulsive decisions that may diverge from their true preferences. An important feature of their framework is that transition into the hot mode is triggered by contextual cues associated with the addictive good. In my setting, the presence of a second pack serves as such a cue, creating the temptation to smoke beyond one's actual need. For a smoker with a low consumption need, having two packs readily available increases the likelihood of entering the hot mode, leading to unintended smoking that would not have occurred had only one pack been purchased.

A key aspect of my model is that the smoker may not have perfect knowledge of the probability of excess consumption π. Instead, when making their purchase decision at $t = 1$, they form beliefs based on a *perceived probability* $\tilde{\pi} \in [0, \pi]$. If the smoker draws a low consumption need c_L, their perceived net present cost of buying two packs is given by

$$\tilde{u}_2 = 2\alpha p_2 + \tilde{\pi}(\delta u_0 + \delta h).$$

The smoker chooses to purchase two packs if and only if $\tilde{u}_2 < u_1$, or equivalently,

$$\alpha p_1(1 + \delta) - 2\alpha p_2 + \delta h > \delta\tilde{\pi}(u_0 + h). \tag{1}$$

On the left-hand side, the first two terms capture the net monetary savings. The third term is the discounted savings in shopping cost. The right-hand side term represents the expected net present cost of smoking an unwanted pack. A smoker is willing to forgo the multipack discount and incur an extra shopping trip if the perceived risk of overconsumption is high enough. Equation 1 has another intuitive implication: smokers who are more *sophisticated*, i.e. those who anticipate their self-control problems and have high $\tilde{\pi}$, are more likely to avoid buying two packs. This relates to a central premise of Bernheim and Rangel (2004): "addicts understand their susceptibility to cue-triggered mistakes and attempt to manage the process with some degree of sophistication." In my context, passing up a discount and accepting the inconvenience of an additional shopping trip serve as a commitment device that helps sophisticated smokers avoid excess consumption.

3 Empirical Setting & Data

3.1 Background on Cigarette Pack Size & Quantity Discounts

In most parts in the world, a standard cigarette pack contains 20 sticks. Many countries have also mandated 20 as the minimum number of cigarettes allowed per pack.[9] The primary reason for this policy is to increase

[9] For example, the United States, Malaysia, Vietnam.

the upfront cost of cigarettes, making them less accessible to teenagers and kids. An additional reason is that a pack of 20 provides sufficient area for the health warning messages to occupy a significant portion of the front and back of the packaging.[10] In places where packs of 10 are still allowed, like India, proposals to increase the minimum pack size are in considerations. In contrast, few places restrict the maximum pack size.[11] This may be because firms have little incentive to offer larger pack sizes when many smokers seek to cut back. Attempts to introduce packs of 25 in the U.S. during the 1990s were largely unsuccessful, in part due to "smokers' stated objectives of wishing to cut down on smoking" (Brown and Williamson 1997).

In the U.S., cigarettes are among the most commonly purchased items at gas station convenience stores and are typically sold in packs of 20. They are almost always stored on shelves directly behind the cashier at checkout, with prices prominently posted. Many stores offer multipack discount deals, such as "buy 2, get $1 off" where customers pay full price for a single pack but receive a discount when buying in multiples of two. These discounts usually apply at the brand level, allowing for mix-and-match of different products within the same brand family. Like prices, these promotions are prominently advertised, often directly on the shelf or even outside the store. Figure 1 shows an example of such a multipack promotion. While some stores offer these deals only through loyalty programs, this paper focuses on multipack deals available to all customers.

Figure 1: Example of a Cigarette Multipack Discount

Source: Photograph by the author.

[10] See this article

[11] A comprehensive overview of global pack size regulations can be found in Blackwell et al. (2020).

3.2 Data Description

The main data for this study is a transaction-level file from over 30,000 U.S gas stations with a convenience store on site.[12] It is provided by a technology company. The data records purchases of all products, including cigarettes, made from July 2019 through December 2024. For most of the analyses, I focus on a subset of 121 stores in Maryland, which experienced two cigarette tax increases during the observed period,[13] creating useful price variations for my structural model. To minimize the impact of the COVID-19 onset, I restrict the analysis to the period from August 2020 onward.

For each cigarette transaction, the data include the store ID, timestamp, product GTIN (Global Trade Item Number) code, number of packs purchased, and the per-pack price. The GTIN code provides information on the brand, flavor, and pack size. The vast majority of transactions (99.6%) involve the standard pack size of 20 cigarettes, so I treat the number of packs as the only relevant dimension of quantity variation. I aggregate products to the brand level and construct a price series by store, week, brand, and number of packs using the modal prices across all transactions, rounded to the nearest cent. This allows me to infer whether a store offers multipack discounts in a given week by comparing the modal 1-pack price to the modal 2-pack price.

The data also allow tracking transactions over time for a subset of consumers, using a combination of the first six and last four digits of their credit card numbers. For the main analyses, I focus on individuals who made at least 10 cigarette purchases and averaged at least 2 packs per month over the observed period. I exclude outlier consumers based on extreme purchasing behavior.[14] The final sample includes 10,607 individuals. Table 1 reports the summary statistics of these smokers.

[12] These stores span 7,861 cities across 48 states and collectively represent approximately 20% of all U.S. gas stations.

[13] The excise tax on cigarettes was increased from $1.75 to $3.75 per pack on March 14, 2021, and raised again to $5.00 on July 1, 2024.

[14] Outliers are defined as individuals who, on average per week, meet any of the following criteria: spend more than $150, make more than 7 store visits, purchase more than 252 ounces of beverages, or buy more than 10 cigarette packs.

Table 1: Summary Statistics of Smokers

	Mean	Median
Packs/week	2.23	1.76
Discount exposure (% trips)	58	67
2-pack share (%)		
With discount	42	33
Without discount	37	25
Buy cigarettes (% trips)	58	58
Conditional on trips with cigarettes purchased		
Buy nothing else (%)	41	39
Share of cigarette spending (%)	51	55
Buy gas (%)	8	3
Buy beverages (%)	16	6
Non-cigarette spending ($)	0.28	0.00

I also examine these individuals' purchasing patterns in beverages and chocolate, two categories with frequent multi-unit discounts at these convenience stores. For beverages, I construct a similar price series by store, week, brand, size (in ounces), and number of units purchased, focusing on standard single-serve sizes usually stored in coolers, specifically 12 to 20 ounces. For chocolate, I construct the price series by store, week, brand, product size, and number of units purchased.

3.3 Data Assumptions

The data have two limitations. First, although I observe individual-level transactions, I do not have access to consumers' complete purchase histories. While gas station convenience stores are a common outlet for cigarette purchases, accounting for 69.1% of sales according to Kruger et al. 2017, consumers may also buy cigarettes elsewhere. To alleviate this concern, I restrict the sample to individuals who buy cigarettes at these stores relatively frequently, at least twice a month. Table 1 shows that, for this group, cigarettes appear to be a primary reason for visiting these stores: cigarettes are purchased on 58% of trips; among cigarette trips, consumers buy no other items 41% of the time; and average spending on non-cigarette items is only $0.28. These patterns suggest that cigarettes are often the sole purpose of the visit. While the concern about missing data remains, this evidence suggests that individuals in my sample likely make a substantial share of their cigarette purchases at these stores.

Secondly, the data only allow me to observe purchase quantities and duration between purchases, not actual cigarette consumption paths. This is a common limitation of retail panel data. To proceed, I make the assumption that smokers buy cigarettes only when they are out of stock. Several patterns in the data support this assumption. First, most purchases are small: 89% of cigarette trips involve buying just one or two packs, and these account for 84% of all cigarette packs sold. This suggests that consumers tend to buy cigarettes for short-term use rather than in bulk. Second, I find no evidence consistent with stockpiling in the data. A hallmark of stockpiling, documented in prior literature, is that consumers return to a store earlier than usual when they anticipate lower prices, suggesting that they still have inventory at home but choose to purchase ahead of need (Erdem et al. 2003, Gordon and Sun 2015, Kim and Ishihara 2025). For cigarettes, however, base prices are generally stable across time and across stores in the absence of tax changes (see Appendix Figure 12). In my sample, 40% of smokers consistently pay the same price, and more than 95% face price fluctuations of less than 50 cents. Furthermore, multipack discounts, if offered, tend to be persistent over time rather than as temporary promotions, which further reduce the incentive to stockpile (see Appendix Figure 13). Lastly, I find no evidence that discounts lead to purchase acceleration: smokers do not appear to return to stores more quickly in response to multipack discounts (see Appendix Table 12). Taken together, these patterns suggest that cigarette purchases in this context are closely aligned with consumption needs.

Since the vast majority of cigarette purchases involve buying one or two packs, I restrict my analysis to these transactions. I also abstract away from brand differentiation, as cigarette is a category characterized by strong brand loyalty. In my data, the average smoker chooses their preferred brand in 89% of purchases, and the median smoker does so 97% of the time. Given these simplifying assumptions, the cigarette purchase decision in my framework reduces to a choice between buying one pack or two packs.

4 Empirical Evidence

This section documents patterns consistent with self-control in cigarette purchases and provides causal evidence of overconsumption. In particular, I use an instrumental variable approach to show that having more packs on hand causally increases smoking.

4.1 Evidence of Self-Control

4.1.1 Multipack Discount Avoidance and Frequent Restocking

Figure 2 shows the distribution of cigarette discount usage across smokers. For each smoker, I compute the share of cigarette trips on which the smoker purchases two packs conditional on being offered a multipack discount on their preferred brand. The bimodal shape is striking: many smokers almost always use discounts, while a similarly large group rarely or never do so.

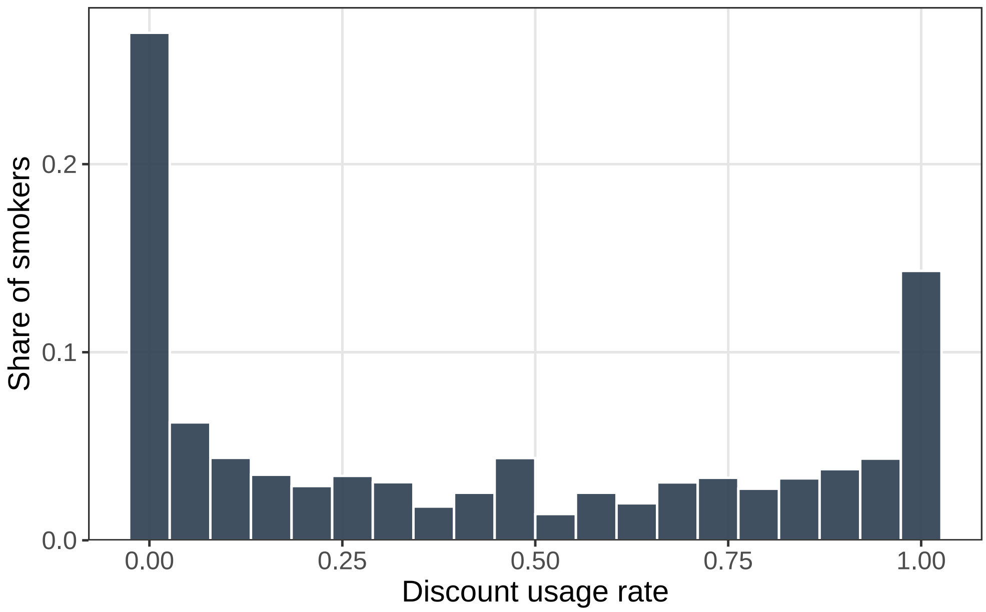

Figure 2: Distribution of Multipack Discount Usage Across Smokers

Note: For each smoker, I compute the probability of purchasing two packs conditional on visiting a store offering a multipack discount. This figure shows the distribution of these probabilities across smokers. 23.7% of smokers always ignore the discounts, whereas 12.6% always utilize them.

The key question, then, is why so many smokers do not take up multipack discounts. As noted earlier, these discounts are prominently displayed at the checkout counter, making it unlikely that smokers overlook them due to inattention. Several other mechanisms could nonetheless account for low discount take-up. Some smokers may smoke infrequently and therefore find two packs excessive. Others may be deterred from paying the higher upfront cost of two packs due to liquidity constraints.[15] It is also possible that some smokers know they can find better prices elsewhere and purchase cigarettes at gas stations only when necessary. Each of these explanations implies a low frequency of single-pack purchases. Thus, if these mechanisms were the primary drivers, a single-pack purchase should typically be followed by a relatively

[15]This explanation is less compelling in my setting because all smokers in the sample use a credit card, which should mitigate short-run liquidity constraints, although it does not fully rule them out.

long interval before the next cigarette trip.

Figure 3 plots the empirical cumulative distribution of the number of days it takes for smokers to restock after buying one pack of cigarettes. The figure therefore represents how long one pack typically lasts.

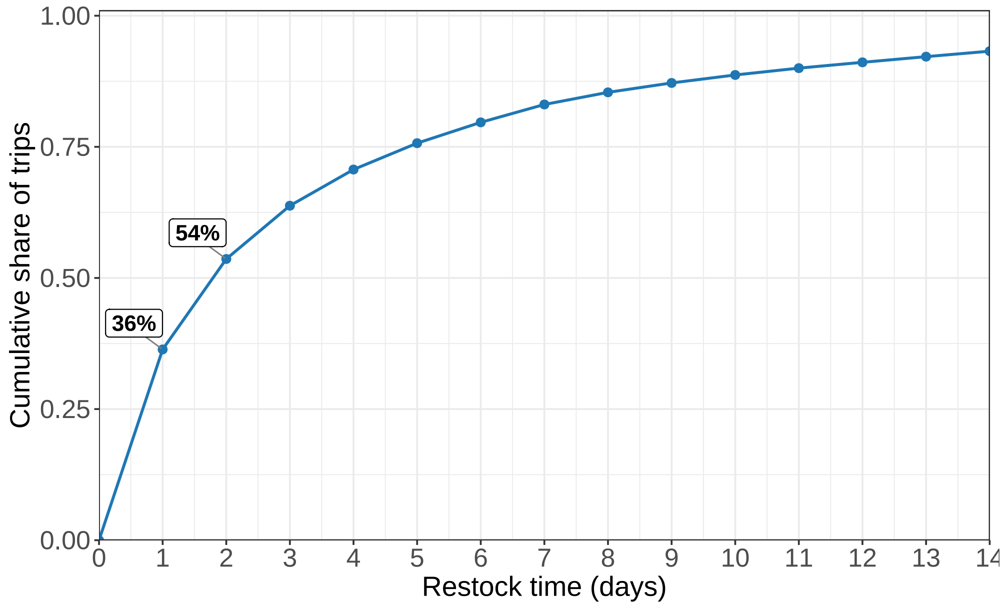

Figure 3: Restock Time Distribution Following Single-Pack Purchases

Note: This figure shows the empirical cumulative distribution of restock times (in days) following trips where consumers purchased a single pack of cigarettes. The average restock time is 4.95 days, the median is 2 days, and the interquartile range is 1-5 days. Approximately 6.8% of such trips are followed by a restock time longer than 14 days, and 1.6% exceed 30 days. In 56% of these trips, a multipack discount is available.

Figure 3 reveals a striking pattern: smokers often return to the store shortly after purchasing a single pack. Over 90% of trips are followed by another visit within two weeks. Notably, 36% of single-pack purchasers return the very next day, and more than half (54%) restock within two days. The 75th percentile is just five days. Moreover, a quantity discount is available in over half of these trips (56%). This pattern holds even when restricting to trips where a multipack discount is available; that is, many consumers purchase a single pack despite being eligible for a multipack discount, and then return soon afterward (see Appendix Figure 11). This behavior is puzzling: buying two packs would save both the hassle of an extra trip and some money, yet many smokers consistently opt for more frequent single-pack purchases.

4.1.2 Multipack Discount Take-up in Other Product Categories

While buying one pack and returning soon afterward is consistent with self-control, it could also reflect myopia or low price sensitivity. In other words, smokers may place little value on discount deals, either because they only focus on current consumption or because they perceive the savings as negligible. As shown in Equation 1, even a completely naive smoker with no self-control (i.e., $\tilde{\pi} = 0$) may still choose to buy a single pack and return early if the following inequality holds:

$$\alpha p_1(1 + \delta) + \delta h - 2\alpha p_2 > 0.$$

In this case, even when there is a discount ($p_2 < p_1$), the single-pack purchase may remain more attractive if the smoker heavily discounts the future (low δ), or is relatively insensitive to price (low α). These factors are not specific to cigarettes and should affect responses to discounts more broadly. If they drive the heterogeneity in Figure 2, then smokers who rarely take cigarette multipack discounts should also be less likely to take advantage of multi-unit discounts in non-cigarette products.

To explore this possibility, I examine smokers' responses to multipack discounts in other product categories. A key challenge is finding suitable comparison products. Few gas-station items, like cigarettes, are both commonly purchased and frequently offered with multipack discounts. Beverages come closest to meeting these criteria. I also consider chocolate as an additional comparison category. Although beverages and chocolate may also be unhealthy, self-control concerns are likely more salient for cigarettes, given the stronger addictive properties of nicotine and the prevalence of regulations aimed at limiting their consumption.

A second complication is that, unlike cigarettes, consumers do not have a similarly well-defined preferred product in beverages or chocolate. Smokers tend to be highly brand loyal, allowing me to construct a single discount indicator based on whether a smoker's preferred cigarette brand is offered with a multipack discount. By contrast, consumers' choices in beverages and chocolate are more diffuse, and multipack discounts typically apply only to specific brands and package sizes. Because we do not observe which products a consumer considers on each trip, it is difficult to define product-level discount exposure in a meaningful way. I therefore construct a category-level indicator for whether any product in the category is offered with a multipack discount.[16] I then measure consumers' responses using the realized savings from their purchases.[17]

I classify smokers into two groups based on cigarette discount take-up, using a 50% cutoff. Smokers in the low cigarette-discount group take the cigarette multipack discount on fewer than half of eligible cigarette trips, while smokers in the high cigarette-discount group take it on at least half of such trips. If cigarette discount avoidance were primarily driven by low price sensitivity or myopia, then smokers in the low cigarette-discount group should also save less from multipack discounts in other categories. Table 2 shows otherwise.

[16] A product is defined at the brand–package level (e.g., Red Bull 12oz, Hershey's 2.55oz).

[17] Realized savings is the difference between the amount paid and the pre-discount amount.

By construction, the difference in cigarette savings is large: smokers in the high cigarette-discount group save \$0.75 per cigarette transaction, compared with \$0.09 for smokers in the low cigarette-discount group. In contrast, the corresponding differences in other categories are economically small. Beverage savings are \$0.12 for the high cigarette-discount group and \$0.11 for the low cigarette-discount group; chocolate savings are \$0.02 and \$0.01. The same pattern holds in percentage terms.

Table 2: Savings from Multipack Discounts by Cigarette Discount Use

	Cigarette Discount Use		
	Low (<50%)	High (≥50%)	Diff.
Beverages	\$0.11	\$0.12	\$0.01
	[1.84%]	[2.18%]	[0.39%]
Chocolate	\$0.04	\$0.05	\$0.01
	[1.05%]	[1.08%]	[0.03%]
Cigarettes (benchmark)	\$0.09	\$0.77	\$0.68
	[0.43%]	[3.58%]	[3.14%]

Notes: Low and high groups are defined by each consumer's cigarette multipack discount take-up rate, measured as the share of cigarette trips on which the consumer takes the multipack discount when offered. Dollar savings are reported in the main row for each category; percentage savings are reported in brackets below. Diff. = High − Low.

In Appendix Section 7, I repeat this comparison in a regression framework that controls for consumer, store, week, and day-of-week fixed effects. Smokers in the high cigarette-discount group still do not realize greater savings from discounts on beverages or chocolate than those in the low group. Taken together, these findings suggest that discount avoidance is specific to cigarettes rather than reflecting a general tendency to avoid discount deals.

4.1.3 Alternative Explanations

Several mechanisms other than self-control could also generate the empirical patterns shown in Figures 2–3. One possibility is the inconvenience of carrying two packs, though this seems unlikely to be the main driver: over 90% of consumers in the sample own a car,[18] cigarette packs are compact and lightweight, and smokers who avoid cigarette discounts still redeem beverage multi-unit discounts at normal rates, even though carrying an extra beverage is arguably more cumbersome than carrying an extra pack of cigarettes.

[18] I infer car ownership from whether the consumer purchased gasoline.

Other explanations include trip-level mental accounting, social-image concerns, and habit. Smokers may avoid an additional pack to stay within a perceived budget, to avoid signaling heavy smoking, or simply because they have developed a routine of buying one pack per trip.

Although the transaction data do not directly reveal smokers' motives, the observed patterns are consistent with extant survey and lab-based evidence showing that many smokers want to cut down and intentionally ration cigarette purchases to limit their own consumption (Brown & Williamson 1997; Centers for Disease Control and Prevention 2024; Marti & Sindelar 2015). The key empirical question is whether such rationing has a meaningful self-control incentive: does having more cigarettes on hand cause smokers to smoke more? If cigarette consumption were exogenous to purchase quantity, as it might be for household staples such as detergent or ketchup, then smokers would have little incentive to ration cigarette purchases for self-control reasons. The next section therefore tests whether cigarette consumption is endogenous to pack size.

4.2 A Test for Overconsumption

4.2.1 Suggestive Data Patterns

First, we examine what happens when smokers buy two packs. Figure 4 plots the empirical cumulative distribution of the number of days it takes for smokers to restock, separately for one-pack and two-pack purchases.

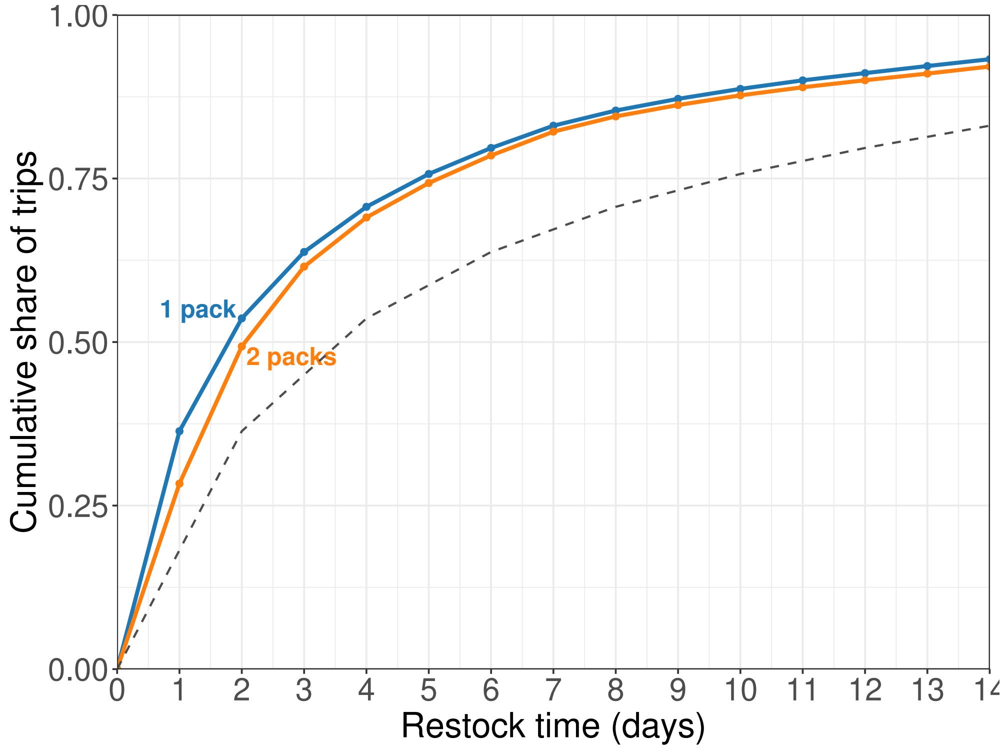

Figure 4: Restock Time Distribution

Note: This figure shows the empirical cumulative distribution of restock times (in days), grouped by whether consumers purchased one pack (blue curve) or two packs (orange curve) of cigarettes. For two-pack trips, the average restock time is 5.56 days, the median is 3 days, and the interquartile range is 1–6 days.

The two curves nearly coincide, indicating that smokers return to the store just as quickly after buying two packs as they do after buying one. This pattern provides suggestive evidence that having more cigarettes on hand may lead to higher consumption. However, the figure may simply reflect self-selection rather than a causal effect. Individuals who tend to smoke more may be more likely to buy two packs in the first place. Even for the same individual, time-varying cigarette demand could jointly determine both the purchase quantity and the time until the next trip. As outlined in the conceptual model in Section 2.2, a smoker might experience high cravings in a given period ($c = c_H$), leading them to purchase more and return sooner.

To assess whether the number of packs itself affects consumption, I focus on identifying two-pack purchases that are likely *intended* to save—analogous to the scenario in the model where the smoker has a low consumption need $c = c_L$ but still chooses to buy two packs. To disentangle the causal effect of having more packs from the effect of underlying consumption demand, I adopt an instrumental variable approach. The excluded instrument is an indicator for whether the visited store offers a multipack discount for the smoker's preferred brand. My empirical strategy tests whether the discount-induced increase in purchase quantity leads to a shorter-than-expected restock time. The next section provides a detailed description of the IV analysis.

4.2.2 Econometric Specification

Let i index individual and t index trip. I examine the effect of purchasing two packs on log$(\tau_{i,t+1})$, where $\tau_{i,t+1}$ is the number of days between trips t and $t+1$ (i.e. the restock time). Denote $q_{it} \in \{1,2\}$ as the number of packs purchased during trip t, the main regression specification is

$$\log\left(\tau_{i,t+1}\right) = \mu_i + \gamma \mathbb{1}[q_{it} = 2] + \mathbf{X}_{it}'\beta + \epsilon_{it}, \tag{2}$$

where μ_i is individual fixed effects, $\mathbf{X}_{it}$ is a set of control variables which includes: logarithms of average consumption over the last 7, 14, and 30 days, store fixed effects, week fixed effects, day-of-week fixed effects, and single-pack price. The coefficient of interest is γ, which measures the degree of overconsumption. Under no overconsumption, γ should equal log$(2) \approx 0.69$, implying that smokers take twice as long to restock when buying two packs as when buying one. An estimate of $\gamma < 0.69$ indicates overconsumption.

The individual fixed effects account for baseline differences in smoking rates and purchase behavior across smokers. Given the addictive nature of cigarettes, recent consumption is a strong predictor of current cravings (Becker and Murphy 1988, Tuchman 2019, Levine 2025). Thus, the recent consumption measures in the covariates help absorb part of this variation. However, there may still be residual variation in current cravings, potentially driving both smoking rate and the number of packs purchased.

To fully address the endogeneity concern, I instrument for q_{it} with D_{it}, an indicator for whether smoker i visits a store offering a multipack deal for i's favorite brand at trip t.

4.2.3 Multipack Discount as an Instrument

Multipack discounts are typically established through contractual agreements between cigarette manufacturers and retailers and are often described as "pervasive and long-standing" (Feighery et al. 2003, Reimold et al. 2023). These brand-level discounts are required to be prominently displayed at the checkout counter. Table 3 reports the summary statistics of cigarette prices and multipack discounts for the three largest brands: Marlboro, Newport, and Camel.

Table 3: Summary Statistics by Cigarette Brand (Average)

	Marlboro	**Newport**	**Camel**
Market share (%)	45	29	5
Price/pack ($)	10.88	11.24	10.74
Price diff. (disc. vs. non-disc. store)	0.08	-0.02	0.03
Discount characteristics			
Frequency (% store/week)	60	42	35
Depth ($)	0.92	1.00	0.70
Length (weeks)	27	13	11

Table 3 confirms that multipack discounts are persistent store-level promotions: Marlboro discounts last on average nearly six months, while Newport and Camel discounts persist for about 13 and 11 weeks, respectively. Hence, the instrument's identifying variation comes from smokers being observed under different store-level discount spells over the 4-year period, both within and across the stores they visit.

The exclusion restriction requires that multipack discount exposure affect restock time only through the number of packs purchased. One concern is supply-side targeting: manufacturers may target stores or neighborhoods with high concentrations of smokers when offering these discounts. However, store fixed effects absorb permanent differences across stores, and smoker fixed effects absorb persistent differences across consumers. Since discounts are long-lasting store-level promotions rather than short-run coupons or personalized offers, they are unlikely to respond to a smoker's temporary cravings on a given trip.

A second concern is demand-side sorting. Since multipack discounts are persistent, smokers may know which stores offer these deals and factor that into their store choice. For example, a smoker experiencing intense cravings on a given day might anticipate buying two packs and deliberately visit a store that offers a multipack discount.

To empirically assess whether this is a concern, I leverage the addictive nature of cigarettes: cravings build up over time and are therefore partially inferred from recent consumption. In other words, cravings have an observable component. If store choice is systematically driven by unobserved cravings, we should see that smokers with higher recent consumption are more likely to visit stores offering multipack discounts. To test this, I run the following regression

$$D_{it} = \nu_i + \psi \log(\hat{c}_{it}) + \tilde{\mathbf{X}}_{it}'\theta + \xi_{it}, \tag{3}$$

where ν_i is individual fixed effects, $\tilde{\mathbf{X}}_{it}$ includes smoker, week and day-of-week fixed effects, and $\hat{c}_{it}$ is a measure of recent cigarette consumption. I use two alternative measures for $\hat{c}_{it}$: (i) $\bar{c}_{it}^{7}$, average daily consumption over the past 7 days, and (ii) $\bar{c}_{it}^{14}$, average daily consumption over the past 14 days. If $\psi > 0$, this would suggest that smokers with higher recent consumption are systematically selecting into stores with discounts, potentially violating the exclusion restriction.

The results are reported in Appendix Table 14. For both the 7-day and 14-day average consumption rate measures, the coefficients are positive and statistically significant at the .001 level. However, their magnitudes are not economically meaningful. In particular, the estimates imply that doubling the recent consumption rate would increase the probability of visiting a store with multipack discounts by only 0.6–0.7 percentage points. This corresponds to a percentage increase of just about 1.2%, a negligible effect. These findings alleviate concerns that short-term spikes in cravings systematically influence store choice.

4.2.4 Evidence of Overconsumption

We report our main IV results in Table 4.

Table 4: Effects of Pack Size on Return Time and Smoking Rate

	IV		
	First stage	Second stage	OLS
$\mathbb{1}$[multipack discount]	0.044***		
	(0.004)		
$\mathbb{1}$[buy 2 packs]		0.391***	0.230***
		(0.085)	(0.004)
Increase in restock time		48%	26%
Increase in daily smoking		35%	59%
Baseline dependent variable	1.384	5 days	5 days
F-statistic	124.1		
Observations	835,196	835,196	835,196

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.
Note: This table reports results from the regression in Equation 2. The unit of analysis is individual-trip. The first column reports the first-stage regression, using an indicator for multipack discount as an instrument for purchasing two packs. The second column reports the corresponding 2SLS estimate. The final column presents the OLS estimate for comparison. All specifications include smoker, store, week, and day-of-week fixed effects, measures of recent consumption, and single-pack price. Standard errors are clustered at the individual level and reported in parentheses.
From the 2SLS coefficients, the increase in return time is calculated as $\exp(0.391) - 1 \approx 48\%$. This implies a 35% increase in daily cigarette consumption, calculated as $2/\exp(0.391)$. The corresponding OLS values are calculated in the same manner.

The first column of Table 4 reports the first-stage results. The high F-statistic of 124.1 indicates that the instrument is relevant. The second column reports the second stage results. The IV estimate of $\hat{\gamma}$ is 0.391, significantly lower than the benchmark value of 0.69 ($p < 0.001$), providing strong evidence of overconsumption. This estimate implies that when smokers have twice as many cigarettes, they take only 48% longer to restock, corresponding to a 35% increase in daily smoking rate. For comparison, the third column reports the OLS estimate. Ignoring the endogeneity of purchase quantity leads to an estimated 59% increase in smoking rate.

Interpretation of IV Results

One important caveat in interpreting the results in Table 4 is the possibility of sharing behavior. Smokers who purchase two packs may be more inclined to share their cigarettes with other fellow smokers. In which case, the observed earlier-than-predicted restock time would not necessarily indicate overconsumption but rather reflect consumption by multiple individuals. As with the underlying consumption decision, sharing behavior is unobserved in the data, and thus this alternative explanation cannot be fully ruled out.

A second concern is that the IV estimate may overstate the degree of overconsumption. Under heterogeneous treatment effects and the usual monotonicity assumption (i.e., a smoker who would otherwise buy two packs is not discouraged by the discount), the IV coefficient captures the effect of buying two packs on smoking rate for *compliers*—smokers who purchase two packs only because a discount is available. Discount sensitivity could be correlated with an inherent tendency to smoke; for example, individuals under stress may be more likely both to respond to discounts and to increase consumption. In this case, the IV estimate could be upward biased.

In my setting, such correlation is unlikely for two reasons. First, it is unclear that the discount meaningfully differentiates high-intensity smokers: individuals who are already smoking heavily or under stress are likely "always-takers", buying two packs regardless of the discount. Second, I test this directly by re-estimating the first stage of Equation 2 with an additional interaction term between discount availability D_{it} and log$(\hat{c}_{it})$, where log$(\hat{c}_{it})$ is again a measure of recent consumption and proxies for factors driving current cravings (including stress). The interaction coefficient (see Appendix Table 15) is small: a smoker with five times higher recent consumption is only 1.3 percentage points more likely to take up the discount. While I cannot rule out bias arising from compliers whose discount sensitivity is correlated with current cravings, the evidence suggests such heterogeneity is limited, implying little scope for bias from this channel.

Another concern is a potential price response: discounts may enable some smokers who wish to consume two packs to afford them. On average, the discounts reduce the price by about 5%. Given a price elasticity of cigarettes in the sample of -0.7 (see Appendix Table 16), this implies an increase in consumption of only about 3.5%—even ignoring stockpiling behavior and the fact that quantity discounts apply only to purchases of two or more packs. This is a small fraction of the observed 35% increase. Moreover, as shown in Appendix Figure 14 and Appendix Table 13, conditional on buying two packs, restock times are identical regardless of whether those packs were purchased on discount or not. Thus, any price response is likely to account for only a negligible share of the observed increase in consumption.

The finding that having more packs leads to increased smoking implies that smokers have a clear incentive to avoid buying multiple packs, but only under two conditions. First, they must have a desire to reduce their smoking. This is a reasonable assumption, given the well-documented health risks of smoking and consistent survey evidence showing that many smokers express a desire to quit or cut back. Second, they must recognize that purchasing more packs can lead to increased consumption. While most smokers may understand the general harms of smoking, they may not anticipate that buying in larger quantities could trigger greater use. Therefore, only smokers who are sophisticated about the risk of overconsumption will choose to forgo multipack discounts. In the following sections, I formalize this idea within a structural framework that quantifies heterogeneity in self-control and simulates policy counterfactuals.

5 A Structural Choice Model

In this section, I specify a structural model to estimate the extent of overconsumption and self-control awareness in the sample. Previous models of cigarette consumption using individual-level data typically consider purchases aggregated at the weekly level or longer (Gordon and Sun 2015, Tuchman 2019, Levine 2025). However, such aggregation smooths out important variations in purchase quantities across trips. For example, a weekly total of two packs could result from either two separate single-pack purchases or one two-pack purchase—a crucial distinction for identifying self-control sophistication. To preserve this variation, I model cigarette purchase decisions at the trip level.

The model builds on the theoretical framework introduced in Section 2.2. Recall that the key feature of this framework is the distinction between the true probability of overconsumption, denoted by π, and the smoker's perceived probability of overconsumption, denoted by $\tilde{\pi} \in [0, \pi]$. A smoker's level of sophistication is defined by how closely their perceived risk $\tilde{\pi}$ aligns with the true risk π. Substantively, as implied by Equation 1, this manifests in how often a smoker opts for frequent single-pack purchases over multipack discounts.

To capture heterogeneity in sophistication, I classify smokers into four behavioral types:

1. **Casual:** $\tilde{\pi} = \pi = 0$. These smokers have no overconsumption problem and consume only what they

need, regardless of how many packs they purchase.

2. **Sophisticated:** $\tilde{\pi} = \pi > 0$. These smokers are prone to overconsumption but are fully aware of this tendency and adjust their purchase behavior accordingly.

3. **Naive:** $\pi > 0$, $\tilde{\pi} = 0$. Like the sophisticated, these smokers overconsume when buying multiple packs, but they mistakenly believe they are casual.

4. **Partially naive:** $0 < \tilde{\pi} < \pi$. These smokers recognize their tendency to overconsume, but underestimate it.

5.1 Model Specifications

As before, let i index individuals and t index trips. The smoker makes a joint decision: how many packs to purchase, denoted by q_{it}, and how much to consume per day, denoted by c_{it}. The smoker will run out of stock and have to return to the store after $\tau_{it} = q_{it}/c_{it}$ days. I conduct my analysis for the case $q_{it} \in \{1, 2\}$.

The objective of the model is to capture heterogeneity in self-control awareness. Following Miller et al. (2025), I assume that smokers differ in their true and perceived probabilities of overconsumption, denoted by π_i and $\tilde{\pi}_i$. In the current specification, I allow for a latent casual type with $\pi_i = 0$. Among non-casual smokers, I assume a common true probability of overconsumption π, but allow perceived overconsumption risk $\tilde{\pi}_i$ to vary across latent types.

A crucial element of the model is the relationship between q_{it} and c_{it}. Following purchase-consumption models typically used in the stockpiling literature (Ching & Osborne 2020; Erdem et al. 2003), I assume that before each trip, the smoker draws a latent consumption need $\tilde{c}_{it}$, and then chooses how many packs to buy, denoted by q_{it}, to satisfy this need. Notice the tilde notation, which highlights a key distinction in my model: realized consumption c_{it} may differ from latent need $\tilde{c}_{it}$. In particular, if the smoker purchases two packs, then there is a probability π_i that they overconsume, that is, $c_{it} > \tilde{c}_{it}$.

I assume the consumption need $\tilde{c}_{it}$ is determined stochastically based on a state variable k_{it}, which summarizes recent consumption history. Motivated by the quadratic consumption utility specifications commonly used in rational addiction models (Becker & Murphy 1988; Gordon & Sun 2015; Kim & Ishihara 2025), I parameterize the distribution of consumption need as:

$$\rho(c_\ell \mid k_{it}; \Theta_c) = \frac{\text{exp}\{\gamma_1 c_\ell + \gamma_2 c_\ell^2 + \gamma_3 c_\ell k_{it}\}}{\sum_{\ell'} \text{exp}\{\gamma_1 c_{\ell'} + \gamma_2 c_{\ell'}^2 + \gamma_3 c_{\ell'} k_{it}\},}$$

where $\Theta_c = \{\gamma_1, \gamma_2, \gamma_3\}$. In rational addiction models (Becker and Murphy 1988; Gordon and Sun 2015; Levine 2025), the state variable k_{it} is typically an addiction stock, evolving according to $k_{it} = (1-\varphi)k_{i,t-1} +$

c_{it}. Here, I do not explicitly model this stock accumulation process; instead, I specify k_{it} as the smoker's average consumption over the previous two weeks, which serves as a reduced-form proxy for their current addiction state, similar in spirit to Tuchman (2019).

I discretize observed consumption rate onto a finite grid and assume latent need lies on the same grid,

$$c_{it}, \tilde{c}_{it} \in \{c_1, c_2, \ldots, c_L\}, \qquad c_1 < \cdots < c_L.$$

In estimation, I use the grid

$$\{0.0625,\ 0.125,\ 0.25,\ 0.5,\ 1,\ 2\}.$$

I assume a simple structure for overconsumption: if a smoker with drawn need $\tilde{c}_{it} = c_\ell$ buys two packs and ends up overconsuming, then realized consumption jumps to the next grid point, $c_{it} = c_{\ell+1}$.[19] This one-step increase avoids excessive parametrization while still capturing the key empirical pattern that having more quantity available may increase consumption.

In the data, I observe the pair (q_{it}, c_{it}), where q_{it} is purchase quantity and c_{it} is realized daily consumption. However, the decision to purchase q_{it} packs is made based on the smoker's latent consumption need $\tilde{c}_{it}$, which may differ from c_{it} if overconsumption occurs. Therefore, to connect the observed data to the underlying decision process, I must map observed (q_{it}, c_{it}) to the unobserved pair $(q_{it}, \tilde{c}_{it})$. Under my model, an observed pair $(q_{it}, c_{it}) = (q, c_\ell)$ implies:

$$\text{if } q = 1: \quad (q_{it}, \tilde{c}_{it}) = (q, c_\ell),$$

$$\text{if } q = 2: \quad (q_{it}, \tilde{c}_{it}) = \begin{cases} (q, c_\ell) & \text{with prob. } 1 - \pi_i, \\ (q, c_{\ell-1}) & \text{with prob. } \pi_i. \end{cases}$$

Having established this mapping, I now turn to modeling the pack-size decision q_{it}, conditional on latent consumption need $\tilde{c}_{it}$. Consistent with the theoretical framework in Section 2.2, this decision involves the following tradeoff: buying one pack entails a higher per-unit price and the cost of an additional shopping trip, whereas buying two packs often offers a lower per-unit price but creates the risk of overconsumption. Note that we are no longer operating within a two-period framework of the theoretical model. To fully capture the smoker's decision process, one would have to specify a dynamic model with infinite horizon. Each purchase decision includes not only the current monetary cost but also the expected continuation value of future restocking. Capturing this continuation value would involve additional assumptions about the evolution of consumption and addiction stock, as in Gordon and Sun (2015). To avoid the complexity of such a model

[19] At the upper boundary, $c_{L+1} = c_L$. At the lower boundary, the overconsumption term involving $c_{\ell-1}$ is omitted.

while retaining the key features of overconsumption and self-control, I do not specify a full infinite-horizon dynamic model. Instead, I assume that, at the time of purchase, the smoker considers how to acquire the next two packs. Under this framing, the smoker compares two options: buying two packs now, or making two sequential one-pack purchases, with the second occurring after the first pack runs out. This captures the idea that smokers face a forward-looking self-control problem but may reduce it to a static decision between “buy two now” versus “buy one now and one later”. This assumption is supported by two empirical patterns: most purchases involve either one or two packs, and two-pack discounts are prominently featured in marketing, making these options particularly salient in the smoker’s decision-making.

The timing of restocking is determined by the consumption need $\tilde{c}_{it}$. With a consumption need of $\tilde{c}_{it}$ packs per day, a single pack will last $\tau_{it} = 1/\tilde{c}_{it}$ days, which is also the smoker’s anticipated restock time when buying one pack. Since they are contemplating how to acquire the next two packs, this implies a sequence of two one-pack purchases: one now, another after τ_{it} days, followed by a second return after $2\tau_{it}$ days. In contrast, if they purchase two packs up front and avoid overconsumption, they can delay their next trip until $2\tau_{it}$ days.

The utility index of each purchase decision can be written as

$$u_{1it} = \alpha_i p_{1it} \left(1 + \delta^{\tau_{it}}\right) + \delta^{\tau_{it}} h + \delta^{2\tau_{it}} v_{it} + \epsilon_{1it},$$

$$\tilde{u}_{2it} = 2\alpha_i p_{2it} + \xi + \delta^{2\tau_{it}} v_{it} + \tilde{\pi}_i u_0 - (1 - \tilde{\pi}_i)\delta^{\tau_{it}} h + \epsilon_{2it}.$$

Here δ is the daily discount factor, h is the shopping cost, v_{it} is the continuation value, $\tilde{\pi}_i$ is the perceived probability of overconsumption, u_0 is the perceived cost of overconsumption, and $\epsilon_{1it}, \epsilon_{2it}$ are i.i.d. Type 1 Extreme Value errors. The tilde notations for $\tilde{\pi}_i$ and $\tilde{u}_{2it}$ indicate that smokers may misperceive the true probability of overconsumption, and hence the total utility from buying two packs.

I allow the price coefficient to vary with the smoker’s behavior in beverage multipack discount category:

$$\alpha_i = \alpha + \alpha_b \text{bev\_disc}_i,$$

where bev_disc_i is the smoker’s average discount savings from beverages. I also include a common two-pack intercept ξ, which may capture storage costs, inconvenience, or other two-pack preferences not directly related to overconsumption.

The term $\delta^{2\tau_{it}} v_{it}$ appears in both equations and thus can be normalized to zero without loss of generality. This reduces the consumer’s quantity choice to a static decision. As a technical assumption, I impose that the consumer always chooses two packs when $\tilde{c}_{it} = 2$.[20]

Let $\mathbf{p} = (p_1, p_2)$ be the price vector and $\Theta_p = \{\alpha, \alpha_b, \delta, h, \xi, \pi, u_0\}$ be the structural parameters entering the purchase decision. Denote

$$V(\tilde{c}, \mathbf{p}; \Theta_p, \tilde{\pi}) = 2\alpha_i p_2 - \alpha_i p_1(1 + \delta^{1/\tilde{c}}) + \xi + \tilde{\pi} u_0 - (1 - \tilde{\pi})\delta^{1/\tilde{c}} h.$$

It follows that the probability that consumer i chooses to buy two packs during trip t, conditional on consumption need $\tilde{c}_{it}$ and prices $\mathbf{p}_{it}$, is

$$\Pr\left(q_{it} = 2 \mid \tilde{c}_{it}, \mathbf{p}_{it}; \Theta_p, \tilde{\pi}_i\right) = \begin{cases} 1 & \text{if } \tilde{c}_{it} = 2, \\ \dfrac{\exp[V(\tilde{c}_{it}, \mathbf{p}_{it}; \Theta_p, \tilde{\pi}_i)]}{1 + \exp[V(\tilde{c}_{it}, \mathbf{p}_{it}; \Theta_p, \tilde{\pi}_i)]} & \text{otherwise.} \end{cases} \tag{4}$$

5.2 Estimation

I estimate the model using the joint likelihood of purchase quantities and realized consumption. For each smoker, the likelihood integrates over the four behavioral types defined above, with type probabilities estimated jointly with the structural parameters. Let p_c, p_s, p_n, and p_p denote the shares of casual, sophisticated, naive, and partially naive smokers.

Let $f(\tilde{c}, \mathbf{p}; \Theta, \tilde{\pi}_i)$ denote the probability that consumer i chooses two packs, given their current consumption need $\tilde{c}$, price vector $\mathbf{p}$, perceived overconsumption risk $\tilde{\pi}_i$, and other parameters, as defined in Equation 4. The likelihood of observing $(q_{it}, c_{it}) = (q, c_\ell)$, conditional on type $(\pi_i, \tilde{\pi}_i)$, observed state k_{it}, and prices $\mathbf{p}_{it}$, is:

$$\begin{aligned} \text{if } q = 1: \quad & L(q, c_\ell \mid k, \mathbf{p}; \pi_i, \tilde{\pi}_i) = \rho(c_\ell \mid k; \gamma)\left[1 - f(c_\ell, \mathbf{p}; \Theta, \tilde{\pi}_i)\right], \\ \text{if } q = 2: \quad & L(q, c_\ell \mid k, \mathbf{p}; \pi_i, \tilde{\pi}_i) = \rho(c_\ell \mid k; \gamma) f(c_\ell, \mathbf{p}; \Theta, \tilde{\pi}_i)(1 - \pi_i) \\ & \qquad\qquad + \rho(c_{\ell-1} \mid k; \gamma) f(c_{\ell-1}, \mathbf{p}; \Theta, \tilde{\pi}_i)\pi_i. \end{aligned}$$

When $q = 2$, the first term corresponds to cases where the smoker consumes as planned, while the second term captures cases where the smoker had a lower latent need but ends up overconsuming. Given the

[20]Because a day is the most granular level of observation, a smoker cannot satisfy a consumption need of two packs per day if they purchase only one pack.

trip-level likelihood $L_{it}(\cdot)$, the smoker-level likelihood integrates over the four latent types:

$$\begin{aligned}\mathcal{L}_i(\Theta) = p_c \prod_t & L(q_{it}, c_{it} \mid k_{it}, \mathbf{p}_{it}; 0, 0) \\ &+ p_s \prod_t L(q_{it}, c_{it} \mid k_{it}, \mathbf{p}_{it}; \pi, \pi) \\ &+ p_n \prod_t L(q_{it}, c_{it} \mid k_{it}, \mathbf{p}_{it}; \pi, 0) \\ &+ p_p \prod_t L(q_{it}, c_{it} \mid k_{it}, \mathbf{p}_{it}; \pi, \bar{\pi}),\end{aligned} \tag{5}$$

where $p_p = 1 - p_c - p_s - p_n$. The model is estimated by maximizing

$$\hat{\Theta} = \arg\max_{\Theta} \sum_i \log \mathcal{L}_i(\Theta),$$

where the parameter vector is

$$\Theta = \{\Theta_c, \Theta_p\} = \{\gamma_1, \gamma_2, \gamma_3, \alpha, \alpha_b, h, u_0, \xi, \delta, \pi, \bar{\pi}, p_c, p_s, p_n\}.$$

As a validation check, I compare the model-implied overconsumption channel to the reduced-form IV evidence. Section 4.2.4 estimates that buying two packs changes log restock time by $\hat{\beta}_{IV} = 0.391$, with standard error 0.085. Under the model's one-step overconsumption structure and the consumption grid design where $c_{\ell+1} = 2c_\ell$, the structural model implies

$$\beta(\pi) = \log\left(\frac{2}{1+\pi}\right).$$

This follows because, after buying two packs, the smoker consumes at the planned rate with probability $1 - \pi$ and at twice the planned rate with probability π.

After estimating the structural model, I compute

$$\hat{\beta}_{\text{struct}} = \beta(\hat{\pi}) = \log\left(\frac{2}{1+\hat{\pi}}\right)$$

and compare it to the IV estimate. This comparison asks whether the overconsumption channel estimated from the full structural likelihood is consistent with the independent reduced-form evidence on restock timing.

5.3 Identification

This section outlines the identification arguments for the key model parameters. The consumption-need distribution $\rho(c \mid k)$ is anchored by consumption following one-pack purchases. Under the model, when $q_{it} = 1$, realized consumption equals latent need, $c_{it} = \tilde{c}_{it}$. Therefore, the distribution of realized consumption after one-pack trips, and how this distribution varies with recent consumption k_{it}, identifies the parameters $\gamma_1, \gamma_2, \gamma_3$.

The overconsumption probability π is identified from excess upward consumption following two-pack purchases relative to this planned-consumption benchmark. Intuitively, if smokers who buy two packs systematically realize higher consumption than predicted by $\rho(c \mid k)$, the model attributes this excess mass to inventory-induced overconsumption.

The purchase utility parameters are anchored by the latent casual, no-overconsumption benchmark. If casual smokers were observed, their purchase decisions would identify the price sensitivity, discount factor, shopping cost, and common two-pack intercept directly, because the overconsumption terms drop out of their utility. In the current latent-type model, the same logic applies probabilistically: smokers whose histories show little excess consumption after two-pack purchases are more likely to be classified as casual, and their purchase behavior anchors these parameters.

Within this benchmark, the price sensitivity parameters α and α_b are identified from variation in multipack discount depth, together with heterogeneity in beverage discount usage. Given price sensitivity, the discount factor δ is identified from variation in the base one-pack price p_1, especially around the Maryland cigarette tax changes. Conditional on δ, the shopping cost h is identified from how two-pack demand varies with the consumption state c_ℓ: high-consumption smokers run out of a single pack sooner, so avoiding an extra shopping trip is more valuable. The common two-pack intercept ξ is identified from the remaining level of two-pack demand after accounting for price incentives, discounting, and shopping costs.

The remaining type parameters are identified based on how smokers' repeated purchase and consumption histories diverge from the casual benchmark. Casual and naive smokers can have similar purchase behavior, but they differ in consumption transitions: casual smokers do not exhibit excess consumption after two-pack purchases, whereas naive smokers do. Sophisticated smokers are identified by the combination of overconsumption risk and additional avoidance of two-pack purchases. Partially naive smokers fall between naive and sophisticated smokers; their perceived overconsumption risk $\bar{\pi}$ is identified by intermediate two-pack avoidance relative to the sophisticated benchmark.

5.4 Results

The estimated parameters are reported in Table 5- 7.

Table 5: Estimated Parameters of Consumption Need Distribution

Parameter	**Estimate**
$\rho(c \mid k) \propto \exp(\gamma_1 c + \gamma_2 c^2 + \gamma_3 ck)$	
γ_1	2.528
	(0.008)
γ_2	-2.368
	(0.006)
γ_3	3.091
	(0.009)

Table 5 reports the estimates governing the consumption need distribution, $\rho(c \mid k)$. The positive linear term and negative quadratic term mirror the signs typically estimated in quadratic consumption utility specifications (Gordon and Sun 2015; Kim and Ishihara 2025). The positive interaction term indicates that higher recent consumption shifts the distribution toward stronger current cravings, capturing persistence in consumption needs.

Table 6: Estimated Parameters of Purchase Utility

Parameter	**Explanation**	**Estimate**
α	Baseline price coefficient	-0.247
		(0.005)
α_b	Price heterogeneity from beverage discounts	0.048
		(0.016)
h	Shopping cost	-2.220
		(0.021)
u_0	Overconsumption disutility	-8.056
		(0.058)
ξ	Common two-pack intercept	-0.361
		(0.019)
δ	Daily discount factor	0.983
		(0.0003)

The estimates for parameters entering the purchase utility are reported in Table 6. The baseline price coefficient is -0.247, consistent with the IV first-stage results indicating that multipack discounts increase purchase quantity. The coefficient on beverage discount savings, $\alpha_b = 0.047$, implies modest heterogeneity in price sensitivity. This is consistent with the earlier reduced-form evidence that smokers who avoid cigarette discounts still take advantage of beverage discounts, suggesting that their avoidance of cigarette discounts is not simply driven by a general dislike of discounts. The daily discount factor is 0.983, notably lower than the conventional value of 0.9995 often assumed in the literature. The estimate for u_0 implies a substantial disutility from overconsumption, equivalent to a monetary cost of \$32.6, suggesting that smokers experience strong aversion to unplanned smoking.[21] The common two-pack intercept is negative, $\xi = -0.359$, equivalent to about \$1.5 in additional inconvenience or storage cost from buying two packs.[22] Finally, the estimated coefficient for h is -2.219, corresponding to an implied shoppipng cost of \$9.0 per trip.[23] Thus, smokers may still have incentives to buy two packs at once to avoid this cost, even in the absence of multipack discounts.

Table 7: Estimated Parameters of Overconsumption and Type Shares

Parameter	**Explanation**	**Estimate**
π	True probability of overconsumption	0.469
		(0.003)
p_c	Share of casual smokers	0.141
		(0.004)
p_s	Share of sophisticated smokers	0.370
		(0.005)
p_n	Share of fully naive smokers	0.204
		(0.005)
p_p	Share of partially naive smokers	0.285
		–
$\bar{\pi}$	Perceived probability for partially naive smokers	0.214
		(0.001)

Smokers are distributed across all four behavioral types. The model estimates that 14.1% of smokers are casual, meaning they do not face an overconsumption risk. Among the remaining smokers, 37.0% are sophisticated, 20.4% are fully naive, and 28.5% are partially naive. The estimated true probability of

[21] Computed as u_0/η_0.
[22] Computed as ξ/η_0.
[23] Computed as h/η_0.

overconsumption is $\hat{\pi} = 0.469$. Partially naive smokers perceive this risk to be only $\hat{\tilde{\pi}} = 0.214$, or about 46% of the true risk.

Figure 5 illustrates the probability of purchasing two packs for each smoker type. As expected, sophisticated smokers, who are aware of the overconsumption risk, largely avoid multipack purchases, buying a single pack in most trips. Naive smokers do not anticipate overconsumption and therefore display purchase behavior similar to casual smokers; both groups frequently choose two packs. Partially naive smokers fall in between, choosing two packs less often than naive and casual smokers but more often than sophisticated smokers.

As a validation check, I compare the estimated overconsumption probability to the reduced-form IV evidence. To make the two quantities comparable, I transform $\hat{\pi}$ into the implied effect on log restock time. Under the model's one-step overconsumption structure,

$$\hat{\beta}_{\text{struct}} = \log\left(\frac{2}{1+\hat{\pi}}\right).$$

With $\hat{\pi} = 0.469$, this gives

$$\hat{\beta}_{\text{struct}} = \log\left(\frac{2}{1+0.469}\right) = 0.309.$$

The model-implied restock-time effect is somewhat smaller than the IV estimate of 0.391, but the difference is modest relative to the IV standard error of 0.085. In particular, $\hat{\beta}_{\text{struct}}$ lies within the 95% confidence interval of the IV estimate, approximately $[0.224, 0.558]$. Thus, the overconsumption channel estimated from the structural model remains consistent with the reduced-form evidence.

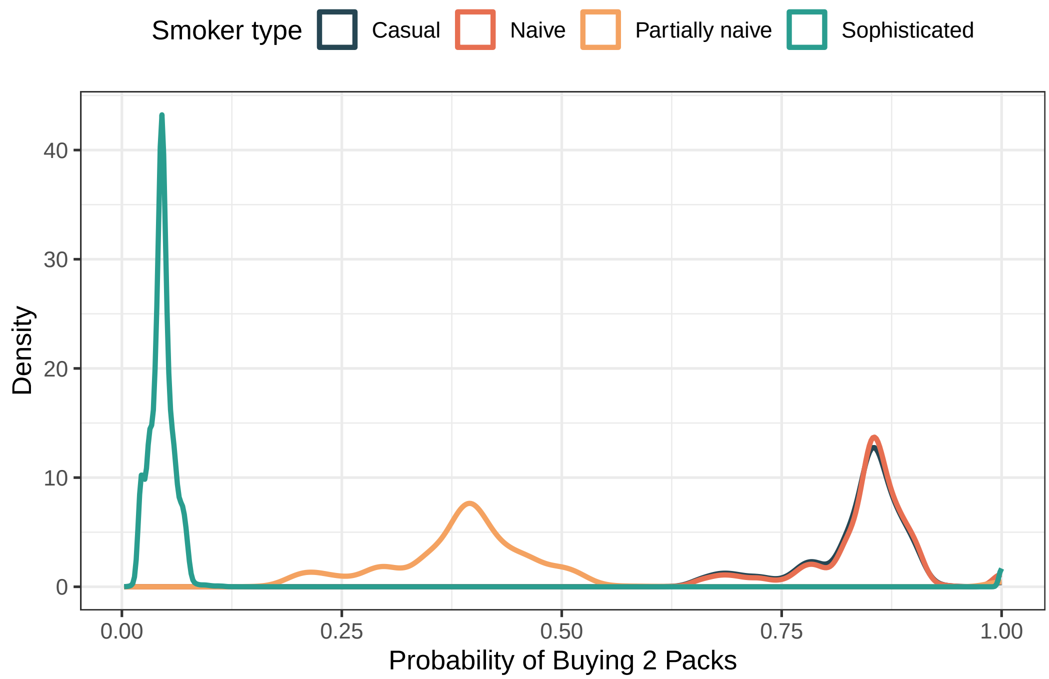

Figure 5: Distribution of Two-Pack Probability by Smoker Type

Note: This figure plots the distribution of the model-implied probability of purchasing two packs for each smoker type. I simulate a panel of 50,000 consumers, each with 100 purchase occasions. Each consumer is assigned to one of the four types as follows: casual with probability equal to the observed share of casual smokers in the data; sophisticated with probability p_s; naive with probability p_n; and partially naive with probability $1 - p_s - p_n$; where p_s and p_n are estimates reported in Table 7. Prices and consumption needs are sampled from the empirical distribution. For each trip, I calculated the probability of purchasing two packs based on Equation 4, using the estimated parameters reported in Table 6. Then, I plot the distribution of such probability across trips, separately for each type.

6 Counterfactual Analyses and Validation

6.1 Setup

In this section, I use the parameter estimates from the structural model to simulate two policy interventions: a ban on multipack discounts and the reintroduction of a half-sized pack. For each counterfactual, I simulate a panel of 100,000 consumers, split evenly between treatment and control groups, each with 100 purchase occasions. Consumers are assigned to one of four latent types using the estimated type shares in Table 7: casual, sophisticated, naive, and partially naive. Prices are sampled from the empirical distribution, while latent consumption needs are drawn from the estimated distribution $\rho(c \mid k)$, with the state variable k initialized from the empirical distribution and updated over the simulated purchase history.

For each purchase occasion, I compute the probability of choosing each available pack size using the estimated purchase-utility parameters from Table 6, and simulate choices accordingly. I then simulate realized consumption. For casual consumers, realized consumption equals latent consumption need. For non-casual consumers, realized consumption equals latent need after one-pack purchases, but increases by one grid level with probability $\hat{\pi} = 0.469$ after two-pack purchases. In the half-pack counterfactual, realized consumption can also fall below planned consumption when consumers choose the half-sized pack.

6.2 Discounts Ban

Cigarette companies allocate the vast majority of their marketing budgets to price discounts (CDC). The IV results in Section 4.2.4 show that these promotions are effective in encouraging smokers to buy more and smoke more, making them a natural target for regulation. In the U.S., the first ban on cigarette discounts was enacted in Providence, RI, in 2012. Since then, cities like Chicago, New York City, and St. Paul have implemented similar policies. At the state level, New Jersey became the first to adopt a ban in March 2020, followed by New York four months later. To date, these remain the only two states with such a ban, and little is known about the effects of these policies.

In this section, I use the structural model estimates to predict the impact of a multipack discount ban on cigarette consumption. In addition, I directly evaluate the effects of such a policy by analyzing the 2020 ban implemented in New York. While there are some caveats, this natural experiment offers a valuable opportunity to assess the real-world impact of the policy and to validate the predictions of my structural model.

6.2.1 Discounts Ban in New York, July 2020

New York banned all cigarette promotions and discounts starting July 1, 2020. To evaluate the effects of this policy, I use weekly store-level data on prices and quantities from July 2019 to July 2021, covering New York and nearby states.[24] The final sample includes 153 stores in New York, 564 in Pennsylvania, 180 in Connecticut, 78 in New Hampshire, 26 in Delaware, and 14 in Rhode Island. I focus exclusively on Marlboro, the only brand with significant multipack activity in New York prior to the ban.[25]

Figure 6 illustrates the effects of the ban on cigarette prices. Prior to the policy, single- and multipack prices in New York and Connecticut followed similar trends. After the ban, the price of multipacks in New York jumped to align with single-pack prices, indicating that multipack discounts became unavailable. Interestingly, there seems to be no change in the trajectory of single-pack prices in New York, suggesting that there was no price response to the ban.

[24] The analysis excludes: (i) New Jersey, which implemented a similar ban four months earlier; (ii) Massachusetts, which enacted a menthol ban in June 2020; (iii) Vermont and Maine, which have limited data coverage; (iv) New York City, which had implemented a discount ban since 2013. I include only stores with at least 8 weeks of observations before and after the ban.

[25] Multipack promotions for Marlboro appeared in 38% of store-weeks, compared to just 0.9% for the next most frequent brand.

Figure 6: Marlboro Prices in New York and Connecticut, July 2019−July 2021

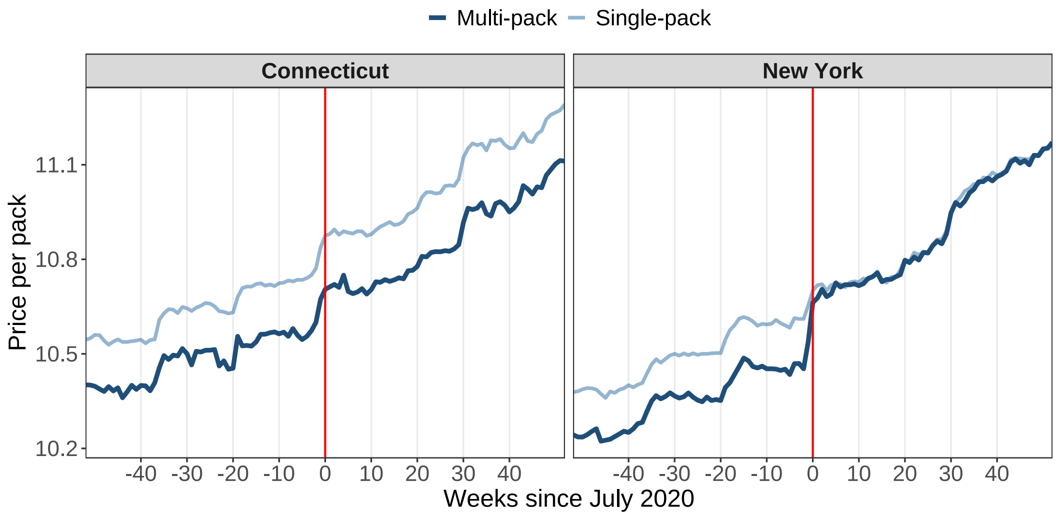

Note: This figure plots weekly Marlboro prices in New York and Connecticut from July 2019 to July 2021. Prices are averaged across stores at the weekly level. The dashed line represents the average single-pack price, and the solid line represents the average multipack price. The red dotted vertical line marks July 2020, when the New York discount ban took effect.

Let s index store and t index week. I run the following regression for each of the four outcomes y_{st}: (1) single-pack price, (2) multipack price, (3) share of two-pack purchases[26], and (4) total packs sold:

$$y_{st} = \alpha_s + \eta_t + \beta(\text{NewYork}_s \times \text{AfterBan}_t) + \epsilon_{st}. \tag{6}$$

The results are reported in Table 8 below. The estimated effects for prices confirm the pattern shown in Figure 6. The ban raised the multipack price by about 13 cents on average, while the single-pack price remained virtually unchanged. Importantly, although the ban reduced the share of two-pack purchases, it had little effect on total cigarette sales, suggesting that the substitution toward one-pack purchases more than offset the reduction. In

[26] defined as the ratio between the number of two-pack purchases and the total number of purchases involving one or two packs

Table 8: Difference-in-Differences Estimates of New York's Discount Ban

	(1) Single-pack price	(2) Two-pack price	(3) Two-pack share	(4) **Total packs**
NewYork$\times$Ban	-0.0004	0.133***	-0.011***	-0.793
	(0.011)	(0.020)	(0.003)	(5.378)
Baseline DV	10.65	10.50	0.20	282.98
Observations	53,925	53,925	53,925	53,925
R^2	0.982	0.969	0.760	0.927

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table reports estimates from a difference-in-differences analysis evaluating the impact of New York's July 2020 ban on cigarette discounts. The estimation window spans 26 weeks before and after the ban. The dependent variables are measured at the store-week level for Marlboro. Control states include Pennsylvania, New Hampshire, Connecticut, Delaware, and Rhode Island. Stores in New York City are excluded due to a preexisting local ban. The sample is limited to stores with at least 8 weeks of observations both before and after the policy change. All regressions include store and week fixed effects. Standard errors are clustered at the store level and reported in parentheses.

6.2.2 Discounts Ban Counterfactual from Structural Model

I use the estimated parameters from the structural model to simulate the effects of a discount ban, following the setup in Section 6.1. Based on earlier evidence of no base price response, I assume that the only difference between treatment and control groups is the availability of multipack discounts. Specifically, consumers in the treatment group face $p_2 = p_1$. I compare the share of two-pack purchases and the consumption rate across groups. While the structural counterfactual is based on individual-level simulations and the policy evaluation uses store-level data, I ensure comparability by computing percentage changes in outcomes. Table 9 reports the simulated effects alongside the estimates from the difference-in-differences analysis.

Table 9: Comparison between Actual Policy Evaluation and Counterfactual Simulations

	% change	
Outcome	DiD	Structural CF
Share of 2-pack purchases	-5.42%	-3.74%
Consumption rate	-0.28%	-0.84%

Note: This table compares the estimated impact of a discount ban using actual difference-in-differences estimates and simulated counterfactuals from the structural model. The DiD estimates use a 26-week window before and after New York's July 2020 discount ban. Simulation setup follows Section 6.1. In the treatment group of the simulation, multipack prices are set equal to single-pack prices. The corresponding 95% confidence intervals for the DiD estimates are approximately $[-8.32\%, -2.52\%]$ for two-pack share and $[-4.00\%, 3.45\%]$ for consumption, both expressed relative to their pre-ban baselines.

The structural model slightly underestimates the decline in the share of two-pack purchases. This discrepancy may stem from differences between New York and Maryland, the shorter time period analyzed, or the fact that the individual-level data used for structural estimation includes only smokers who pay with credit or debit cards, potentially a less price-sensitive group. Nonetheless, both the difference-in-differences analysis and the structural model yield modest decreases in consumption rate (0.28% and 0.84%, respectively), supporting the idea that a discount ban may have limited effects on cigarette consumption.

Figure 7 shows the simulated effects of the discount ban on two-pack purchase probabilities, separately by consumer type. The largest absolute response comes from partially naive smokers, whose two-pack share declines by 3.7 percentage points, or 9.3% relative to baseline. Casual and naive smokers respond only modestly, with two-pack shares falling by about 2 percentage points. Their limited response reflects the fact that, even without discounts, they still have strong incentives to buy two packs to avoid the estimated $9.0 shopping cost. Sophisticated smokers exhibit a larger relative decline of 11.6%, but their baseline two-pack share is only 5.7%, so the absolute change is small.

Figure 8 shows that these purchase-margin responses translate into modest changes in consumption. The largest reduction occurs among partially naive smokers, whose consumption falls by 1.9%, followed by naive smokers at 1.2%. For the other groups, the simulated consumption effects are close to zero.

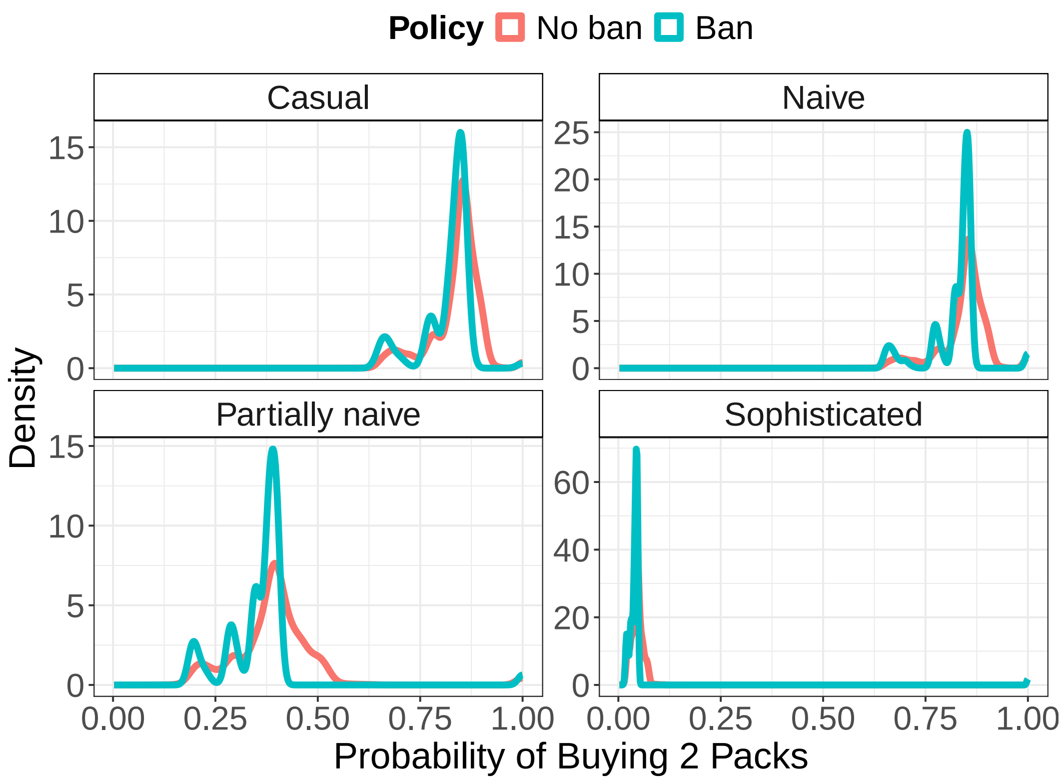

Figure 7: Effects of a Discount Ban on Probability of Buying Two Packs by Smoker Type

Note: These figures compare the estimated impact of a discount ban on the model-implied probability of purchasing two packs, shown separately for each smoker type. Simulation setup follows Section 6.1. In the treatment group, multipack prices are set equal to single-pack prices. For each smoker type and group, I plot the distribution of two-pack purchase probabilities across trips, as implied by Equation 4, with the treatment group in blue and the control group in red.

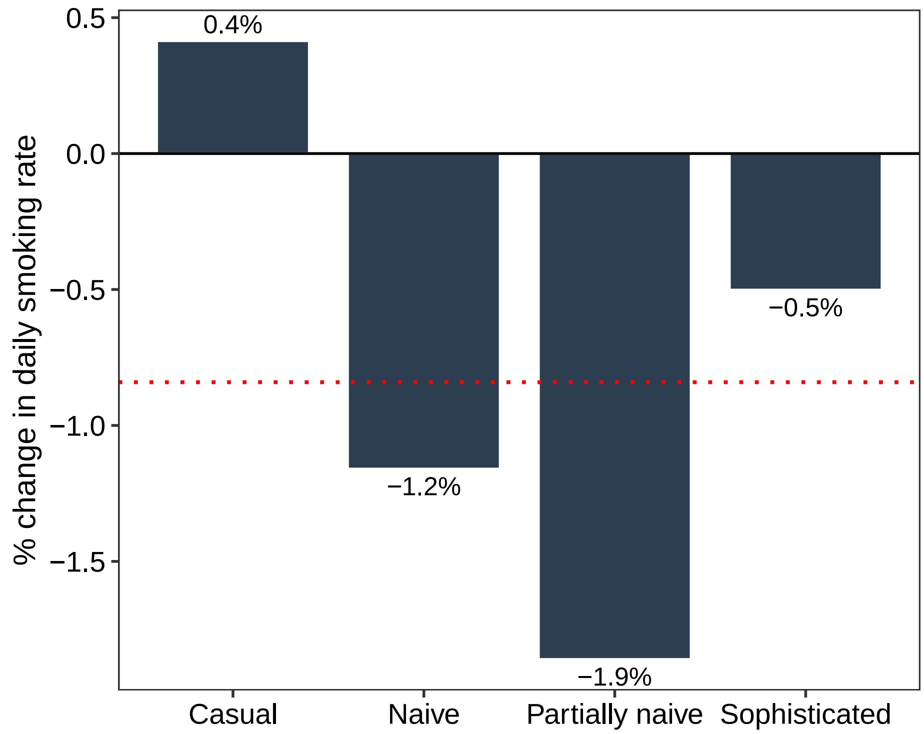

Figure 8: Effects of a Discount Ban on Consumption by Smoker Type

Note: These figures compare the estimated impact of a discount ban on the model-implied consumption, shown separately for each smoker type. Simulation setup follows Section 6.1. In the treatment group, multipack prices are set equal to single-pack prices. The red dotted line indicates the overall effect -0.84%

6.3 Small Pack Reintroduction

In this section, I examine another counterfactual: reintroducing a smaller pack. In 2009, the Family Smoking Prevention and Tobacco Control Act banned the sale of cigarette packs under 20 in the U.S., following similar policies established in other developed countries. While this helps reduce smoking initiation, particularly among youth, a segment of smokers who wish to regulate their consumption might have been negatively affected.

In this counterfactual, I introduce a half-pack with 10 cigarette sticks as a third option for smokers to choose from. The disutilities from buying each pack, conditional on latent consumption need $\tilde{c}$ and per-unit prices $p_{0.5}, p_1, p_2$, are

$$\begin{aligned}
u_{0.5} &= \frac{1}{2}\alpha p_{0.5}\left(1+\delta^{\tau/2}\right) + \delta^{\tau} a p_1 + h\left(\delta^{\tau/2}+\delta^{\tau}\right) - \frac{1}{2}\xi - \frac{1}{2}\tilde{\pi}(u_0 + 2\delta^{\tau/2}h) + \epsilon_{0.5} \\
u_1 &= \alpha p_1(1+\delta^{\tau}) + \delta^{\tau} h + \epsilon_1 \\
u_2 &= 2\alpha p_2 + \xi + \tilde{\pi}(u_0 + \delta^{\tau} h) + \epsilon_2
\end{aligned} \tag{7}$$

where again, $\tau = \frac{1}{\tilde{c}}$, and the continuation value has been normalized to zero. Notice an underlying assumption that after initially purchasing the small pack and returning at $\frac{\tau}{2}$ days for another small pack, the smoker will then return again at τ days to buy a standard pack. This assumption simplifies the modeling of subsequent purchases by fixing the second return to a standard pack. Thus, the monetary cost includes buying one half-size pack now, another after $\frac{\tau}{2}$ days, and a standard pack after τ days. The shopping cost is incurred at both $\frac{\tau}{2}$ and τ. If the smoker buys the small pack, they also "gain" back half the inconvenience cost ξ. For smokers with overconsumption problems, the small pack offers a chance to refrain from returning early at $\frac{\tau}{2}$. I assume the probability of this *underconsumption* is the same as the overconsumption probability π, and that it is perceived similarly by smokers. If the smoker underconsumes, they avoid the shopping cost at $\frac{\tau}{2}$ and gain utility equal to $-\frac{1}{2}u_0$. The intuition here is simple. Consuming an unwanted pack over an interval of $\frac{1}{\tilde{c}}$ incurs a loss of u_0, so refraining from half a pack over the same interval gives back half of the equivalent benefit.

A key determinant of the small pack's take-up and effectiveness is its unit price $p_{0.5}$. In this counterfactual exercise, I consider a range of values for $p_{0.5}$, in particular

$$\frac{p_{0.5}}{2p_1} \in \{60\%, 70\%, 80\%, 90\%, 100\%\}$$

6.3.1 Counterfactual Results from Structural Model

Figure 9 shows take-up of the half-sized pack across different price levels, separately by smoker type. Casual and naive smokers rarely choose the half pack, even when it is relatively cheap, since they perceive little benefit from using it as a commitment device. Partially naive smokers take it up at moderate rates, choosing it about 3%–17% of the time depending on price. Sophisticated smokers respond most strongly: they choose the half pack in about 61% of trips when its total price is 60% of the standard pack price, and still choose it in about 19% of trips even when it costs as much as a standard pack.

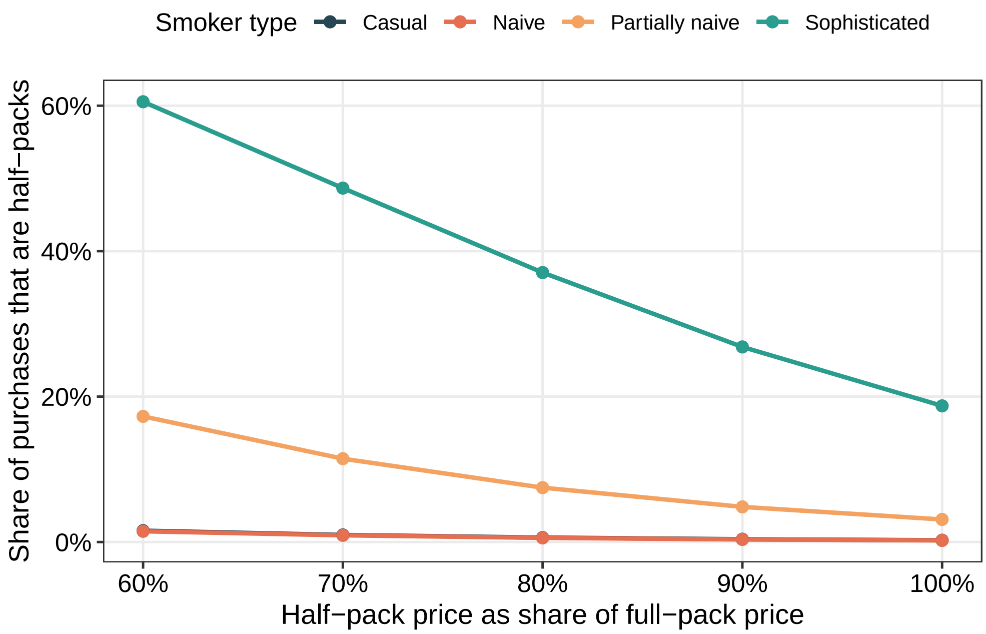

Figure 9: Half-Pack Take-Up by Smoker Type and Price

Note: This figure plots the simulated share of purchases accounted for by the half-sized pack, separately by smoker type. The x-axis varies the total price of the half pack from 60% to 100% of the standard pack price. Simulation setup follows Section 6.1. Purchase probabilities are computed using the multinomial logit model with utilities defined in Equation 7.

Ultimately, the key outcome is total consumption. Figure 10 shows the percentage reduction in smoking from introducing the half-sized pack at different price levels. The largest effects arise among sophisticated smokers, whose consumption falls by about 3%–12%. Partially naive smokers also reduce consumption, especially when the half pack is less expensive. Aggregating across all smoker types, the policy reduces consumption by about 4.5% when the half pack is priced at 60% of the standard pack, and by about 0.8% even when the half pack costs as much as a standard pack. Thus, the half-pack counterfactual can generate consumption reductions comparable to, and often larger than, the simulated discount ban.

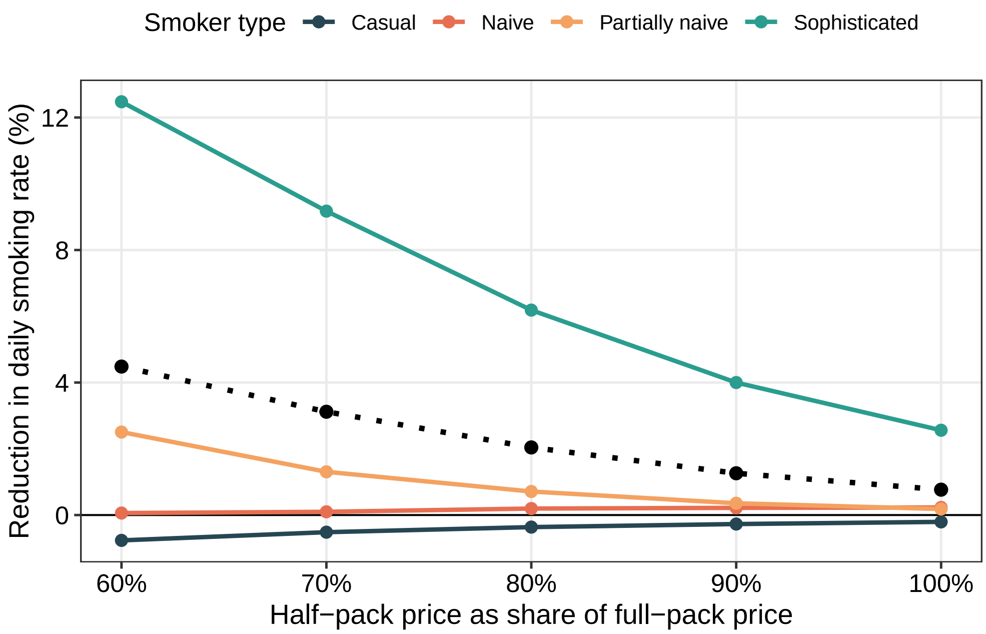

Figure 10: Consumption Reduction from Reintroducing the Half-Sized Pack

Note: This figure plots the simulated percentage reduction in daily cigarette consumption from introducing a half-sized pack. The x-axis varies the total price of the half pack from 60% to 100% of the standard pack price. The dashed line reports the overall reduction across all smoker types. Simulation setup follows Section 6.1.

7 Conclusion

I provide evidence of purchase rationing for harmful goods using real-world transaction data from 10,607 cigarette smokers. A substantial share of smokers consistently avoid quantity discounts and instead make frequent single-pack purchases. This behavior cannot be attributed to price insensitivity or myopia: smokers who forgo cigarette discounts still utilize beverage quantity discounts at similar rates as others. The self-control explanation is further supported by the fact that buying more cigarettes causally leads to more smoking, a result I established through an instrumental variable strategy. Consistent with theories of self-control, extra shopping trips and forgone discounts function as commitment devices to limit consumption.

I estimate the distribution of self-control sophistication using a structural model based on Bernheim and Rangel (2004) and find that most smokers are at least partially aware of their overconsumption risk. These consumers likely benefit from offering smaller pack sizes. Counterfactual simulations confirm this: reintroducing a 10-stick pack reduces overall consumption by about 2% when priced at 80% of the standard pack, and by up to 4.5% when priced at 60% of the standard pack. In contrast, banning quantity discounts yields less than a 1% reduction. These findings shed light on how to design effective interventions to reduce smoking when consumers exercise self-control to a large extent.

This paper can be extended in several directions. First, my analyses largely abstract from the firm's supply-side decisions. While the results offer insight into how pack sizes and prices might affect demand, they do not predict the exact strategy a firm would choose. For example, although I show substantial demand for a half-sized pack priced the same as the standard pack, in reality, if such a product were legally permitted, the firm would likely set a lower price. Another key factor is the actual size of the small pack, which need not be 10 cigarettes. Determining the optimal combination of pack size and price, relative to the common 20-pack, remains an important avenue for future research.

It is important to consider how the insights from this paper might generalize to other vice products beyond cigarettes. While self-control and overconsumption likely exist with other tempting goods such as soda, ice cream, or snacks, detecting these behaviors is more challenging due to their lower purchase frequency. Cigarettes offer a particularly clean setting to study self-control because they are (i) frequently purchased, (ii) easily stored, and (iii) inherently tempting. Few other products share all three properties. For cigarettes, forgoing a multipack discount is a strong signal of self-control, since the smoker will likely need that second pack within a few days. In contrast, someone who forgoes a family-sized pack of chips at a lower unit price may simply prefer not to have chips and choose another snack instead. This is not to say self-control is absent for chips, but empirically, it is far harder to detect in the data.

Another limitation of this paper is the absence of consumer characteristics. The data lack demographic information, which could complement the behavioral insights. While self-control and overconsumption are important, they are intangible and difficult to target directly. Identifying which demographic groups are more likely to be sophisticated, for example, would yield more actionable insights for intervention design. In addition, youth uptake, a crucial consideration for tobacco control policy, is abstracted away here due to data limitations.

References

Ailawadi, K. L., & Neslin, S. A. (1998). The effect of promotion on consumption: Buying more and consuming it faster. *Journal of Marketing Research*, *35*(3), 390–398 (cit. on p. 5).

Allcott, H., Gentzkow, M., & Song, L. (2022). Digital addiction. *American Economic Review*, *112*(7), 2424–2463 (cit. on p. 4).

Alwan, V., Dube, J.-P., & Kong, X. (2024). The identification of the dynamic discrete-choice model with stockpiling. *Available at SSRN* (cit. on p. 5).

Ashraf, N., Karlan, D., & Yin, W. (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *The Quarterly Journal of Economics*, *121*(2), 635–672 (cit. on p. 4).

Bao, Y., Landry, P., & Shi, M. (2025). Forgetful consumers and consumption tracking. *Management Science*, *71*(3), 2311–2331 (cit. on p. 5).

Becker, G. S., & Murphy, K. M. (1988). A theory of rational addiction. *Journal of political Economy*, *96*(4), 675–700 (cit. on pp. 18, 23).

Bernheim, B. D., & Rangel, A. (2004). Addiction and cue-triggered decision processes. *American economic review*, *94*(5), 1558–1590 (cit. on pp. 3, 5–7, 41).

Blackwell, A. K., Lee, I., Scollo, M., Wakefield, M., Munafò, M. R., & Marteau, T. M. (2020). Should cigarette pack sizes be capped? *Addiction*, *115*(5), 802–809 (cit. on p. 8).

Bonfrer, A., Chintagunta, P. K., Roberts, J. H., & Corkindale, D. (2020). Assessing the sales impact of plain packaging regulation for cigarettes: Evidence from australia. *Marketing Science*, *39*(1), 234–252 (cit. on p. 5).

Brown & Williamson. (1997). Analysis of industry pricing dynamics and recommended strategies. *Industry Report*. https://www.industrydocuments.ucsf.edu/docs/#id=njpm0190 (cit. on pp. 8, 16).

Centers for Disease Control and Prevention. (2024, September 17). *Smoking Cessation: Fast Facts*. Retrieved May 5, 2026, from https://www.cdc.gov/tobacco/php/data-statistics/smoking-cessation/index.html (cit. on p. 16).

Chaloupka, F. J., Levy, M. R., & White, J. S. (2019). *Estimating biases in smoking cessation: Evidence from a field experiment* (tech. rep.). National Bureau of Economic Research. (Cit. on p. 6).

Ching, A. T., & Osborne, M. (2020). Identification and estimation of forward-looking behavior: The case of consumer stockpiling. *Marketing Science*, *39*(4), 707–726 (cit. on pp. 5, 23).

Conlin, M., O'Donoghue, T., & Vogelsang, T. J. (2007). Projection bias in catalog orders. *American Economic Review*, *97*(4), 1217–1249 (cit. on p. 5).

Della Vigna, S., & Malmendier, U. (2006). Paying not to go to the gym. *american economic Review*, *96*(3), 694–719 (cit. on p. 4).

DellaVigna, S., List, J. A., & Malmendier, U. (2012). Testing for altruism and social pressure in charitable giving. *The quarterly journal of economics*, *127*(1), 1–56 (cit. on p. 5).

Dobson, P. W., & Gerstner, E. (2010). For a few cents more: Why supersize unhealthy food? *Marketing Science*, *29*(4), 770–778 (cit. on p. 1).

Dubé, J.-P., Luo, X., & Fang, Z. (2017). Self-signaling and prosocial behavior: A cause marketing experiment. *Marketing Science*, *36*(2), 161–186 (cit. on p. 5).

Erdem, T., Imai, S., & Keane, M. P. (2003). Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and economics*, *1*, 5–64 (cit. on pp. 5, 11, 23).

Feighery, E. C., Ribisl, K. M., Clark, P., & Haladjian, H. H. (2003). How tobacco companies ensure prime placement of their advertising and products in stores: Interviews with retailers about tobacco company incentive programmes. *Tobacco Control*, *12*(2), 184–188 (cit. on p. 18).

Giné, X., Karlan, D., & Zinman, J. (2010). Put your money where your butt is: A commitment contract for smoking cessation. *American Economic Journal: Applied Economics*, *2*(4), 213–235 (cit. on p. 4).

Goli, A., Mummalaneni, S., & Chintagunta, P. K. (2024). Making a smooth exit? menthol bans and cigarette sales in massachusetts. *Marketing Science*, *43*(3), 564–589 (cit. on p. 5).

Goli, A., Mummalaneni, S., Chintagunta, P. K., & Dhar, S. K. (2022). Show and sell: Studying the effects of branded cigarette product placement in tv shows on cigarette sales. *Marketing Science*, *41*(6), 1163–1180 (cit. on p. 5).

Gomez, M., & Guevara, M. (1993). Rjr intl. consumer research report. *Islandwide Tracking Study*. https://www.industrydocuments.ucsf.edu/docs/#id=rtpy0013 (cit. on p. 1).

Gordon, B. R., & Sun, B. (2015). A dynamic model of rational addiction: Evaluating cigarette taxes. *Marketing Science*, *34*(3), 452–470 (cit. on pp. 2, 5, 11, 22–24, 29).

Grubb, M. D., & Osborne, M. (2015). Cellular service demand: Biased beliefs, learning, and bill shock. *American Economic Review*, *105*(1), 234–271 (cit. on p. 5).

Gui, G., & Drerup, T. H. (2025). Designing promises with reference-dependent customers: The case of online grocery delivery time. *Available at SSRN 4298782* (cit. on p. 5).

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, *103*(7), 2643–2682 (cit. on p. 5).

Hastings, J. S., & Shapiro, J. M. (2013). Fungibility and consumer choice: Evidence from commodity price shocks. *The quarterly journal of economics*, *128*(4), 1449–1498 (cit. on p. 5).

Hendel, I., & Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, *74*(6), 1637–1673 (cit. on p. 5).

Huang, Y. (2024). Pricing frictions and platform remedies: The case of airbnb. *Available at SSRN 3767103* (cit. on p. 5).

Jain, S. (2012). Marketing of vice goods: A strategic analysis of the package size decision. *Marketing Science*, *31*(1), 36–51 (cit. on p. 1).

Jain, S., & Li, K. J. (2018). Pricing and product design for vice goods: A strategic analysis. *Marketing science*, *37*(4), 592–610 (cit. on p. 1).

Joo, J. (2023). Rational inattention as an empirical framework for discrete choice and consumer-welfare evaluation. *Journal of Marketing Research*, *60*(2), 278–298 (cit. on p. 5).

Kim, J. Y., & Ishihara, M. (2025). Evaluating the effect of soda taxes using a dynamic model of rational addiction. *Forthcoming at Management Science* (cit. on pp. 2, 5, 11, 23, 29).

Kraft, A., & Rao, R. (2024). Behavioral skimming: Theory and evidence from resale markets. *Available at SSRN* (cit. on p. 5).

Kruger, J., Jama, A., Lee, J. G., Kennedy, S., Banks, A., Sharapova, S., & Agaku, I. (2017). Point-of-sale cigarette purchase patterns among us adult smokers—national adult tobacco survey, 2012–2014. *Preventive medicine*, *101*, 38–43 (cit. on p. 10).

Lacetera, N., Pope, D. G., & Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, *102*(5), 2206–2236 (cit. on p. 5).

Levine, J. (2025). Are menthol cigarettes more addictive? a cross-category comparison of habit formation. *Working paper* (cit. on pp. 5, 18, 22, 23).

Lin, T., & Strulov-Shlain, A. (2025). Choice architecture, privacy valuations, and selection bias in consumer data. *Marketing Science* (cit. on p. 5).

Loghmani, E., & Goli, A. (2024). Investigating the impact of advertising on smoking cessation: The role of dtc prescription drug advertising. *Available at SSRN 4775370* (cit. on p. 5).

Marti, J., & Sindelar, J. (2015). Smaller cigarette pack as a commitment to smoke less? insights from behavioral economics. *PLoS One*, *10*(9), e0137520 (cit. on pp. 1, 4, 16).

Milkman, K. L., Minson, J. A., & Volpp, K. G. (2014). Holding the hunger games hostage at the gym: An evaluation of temptation bundling. *Management science*, *60*(2), 283–299 (cit. on p. 4).

Miller, K. M., Sahni, N. S., & Strulov-Shlain, A. (2025). Sophisticated consumers with inertia: Long-term implications from a large-scale field experiment. *Available at SSRN 4065098* (cit. on pp. 5, 23).

Mishra, A., & Mishra, H. (2011). The influence of price discount versus bonus pack on the preference for virtue and vice foods. *Journal of Marketing Research*, *48*(1), 196–206 (cit. on p. 1).

Natan, O. R. (2024). Choice frictions in large assortments. *Marketing Science* (cit. on p. 5).

Reimold, A. E., Lee, J. G., & Ribisl, K. M. (2023). Tobacco company agreements with tobacco retailers for price discounts and prime placement of products and advertising: A scoping review. *Tobacco control*, *32*(5), 635–644 (cit. on p. 18).

Royer, H., Stehr, M., & Sydnor, J. (2015). Incentives, commitments, and habit formation in exercise: Evidence from a field experiment with workers at a fortune-500 company. *American Economic Journal: Applied Economics*, *7*(3), 51–84 (cit. on p. 4).

Schilbach, F. (2019). Alcohol and self-control: A field experiment in india. *American economic review*, *109*(4), 1290–1322 (cit. on p. 4).

Seiler, S., Tuchman, A., & Yao, S. (2021). The impact of soda taxes: Pass-through, tax avoidance, and nutritional effects. *Journal of Marketing Research*, *58*(1), 22–49 (cit. on p. 53).

Strulov-Shlain, A. (2023). More than a penny's worth: Left-digit bias and firm pricing. *Review of Economic Studies*, *90*(5), 2612–2645 (cit. on p. 5).

Sun, B. (2005). Promotion effect on endogenous consumption. *Marketing science*, *24*(3), 430–443 (cit. on p. 5).

Thaler, R. H., & Benartzi, S. (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of political Economy*, *112*(S1), S164–S187 (cit. on p. 4).

Tuchman, A. E. (2019). Advertising and demand for addictive goods: The effects of e-cigarette advertising. *Marketing science*, *38*(6), 994–1022 (cit. on pp. 5, 18, 22, 24).

Wang, Y., Lewis, M., & Singh, V. (2016). The unintended consequences of countermarketing strategies: How particular antismoking measures may shift consumers to more dangerous cigarettes. *Marketing Science*, *35*(1), 55–72 (cit. on p. 5).

Wang, Y., Lewis, M., & Singh, V. (2021). Investigating the effects of excise taxes, public usage restrictions, and antismoking ads across cigarette brands. *Journal of Marketing*, *85*(3), 150–167 (cit. on p. 5).

Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing science*, *17*(4), 317–337 (cit. on pp. 1, 4).

Zhang, S., Chan, T. Y., Luo, X., & Wang, X. (2022). Time-inconsistent preferences and strategic self-control in digital content consumption. *Marketing Science*, *41*(3), 616–636 (cit. on p. 4).

Appendix

Restock Time Distribution Conditional on Being Exposed to Multipack Discounts

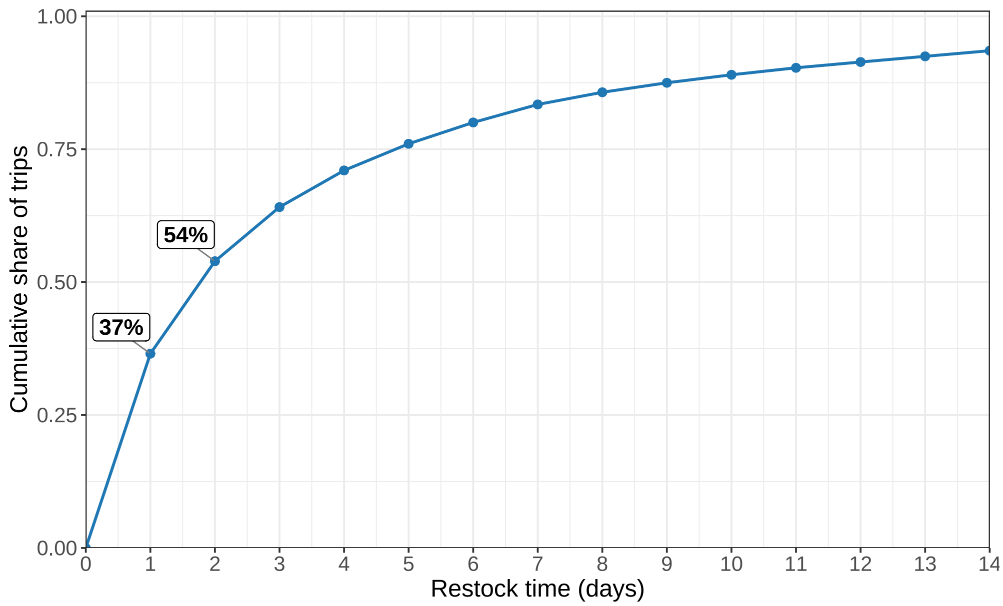

Figure 11: Restock Time Distribution After Rejecting Multipack Discounts

Note: This figure shows the empirical cumulative distribution of restock times (in days) following trips in which consumers purchased a single pack of cigarettes while a multipack deal was available. The average restock time is 4.84 days, the median is 2 days, and the interquartile range is 1–5 days.

Savings from Multipack Discounts of Other Products–Regression with Fixed Effects

Let i index smokers and t index purchase occasions. I estimate the following regression separately for beverages and chocolate:

$$y_{it} = \eta_i + \phi \mathbb{1}[\text{discount}]_{it} + \kappa \left(\mathbb{1}[\text{discount}]_{it} \times \mathbb{1}[\text{High cig. take-up}]_i\right) + \mathbf{X}'_{it}\Delta + \upsilon_{it},$$

where y_{it} denotes discount savings (in dollars or percentage terms), and $\mathbb{1}[\text{discount}]_{it}$ indicates whether the store visited on trip t offers a multipack discount in the relevant category. The coefficient of interest, κ, captures whether smokers in the high cigarette-discount group realize different savings on multipack discounts in beverages or chocolate. The controls $\mathbf{X}_{it}$ include store, week, and day-of-week fixed effects, and η_i denotes smoker fixed effects. As shown in Table 10, the estimated interaction term is close to zero across both categories and outcomes, indicating that consumers with high cigarette discount take-up do not display differential savings from multipack discounts outside cigarettes.

Table 10: Discount Savings by Cigarette Discount Take-up

	DV: savings from multipack discount			
	Beverages		Chocolate	
	(1)	(2)	(3)	(4)
	Dollar	Percent	Dollar	Percent
$\mathbb{1}$[discount]	0.0219***	0.3386***	0.0142***	0.3551***
	(0.0045)	(0.0615)	(0.0033)	(0.0689)
$\mathbb{1}$[discount] $\times$ $\mathbb{1}$[High cig. take-up]	0.000004	0.0099	-0.0011	-0.0296
	(0.0076)	(0.1059)	(0.0058)	(0.1471)
Baseline dependent variable	0.147	2.449	0.023	0.592
Observations	402,852	402,852	42,448	42,448

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.
Note: The unit of analysis is a purchase occasion (beverages in columns (1)–(2), chocolate in columns (3)–(4)), so each sample is conditional on purchasing that category. High cigarette take-up indicates consumers who take cigarette multipack discounts on at least 50% of cigarette trips when such discounts are offered. Standard errors are clustered at the consumer level and reported in parentheses. All specifications include consumer, store, week, and day-of-week fixed effects.

As a further robustness check, I replace the binary discount indicator with a continuous, share-weighted discount index. For each store-week st, I define

$$disc\_index_{st} = \frac{\sum_j share_{jst} \cdot \mathbb{1}[\text{discount on product } j \text{ in } st]}{\sum_j share_{jst}},$$

where $share_{jst}$ is product j's share of purchases in store s and week t. A product is defined by its brand and size (e.g., a 16oz. Coke, 2.55oz Hershey). Thus, $disc\_index_{st}$ is larger when discounts are available on more popular products. I then re-estimate the regression replacing $\mathbb{1}[\text{discount}]_{it}$ with $disc\_index_{it}$. Table 11 shows that the interaction term remains close to zero across both categories and outcomes, indicating that smokers in the high cigarette-discount group do not respond more to multipack discounts in other products, even when the discount environment is stronger.

Table 11: Discount Savings by Cigarette Discount Take-up: Share-Weighted Discount Index

	DV: beverage discount savings		*DV:* chocolate discount savings	
	(1)	(2)	(3)	(4)
	Dollar	Percentage	Dollar	Percentage
$disc\_index$	0.2531***	4.2410***	0.2594***	7.2079***
	(0.0220)	(0.3402)	(0.0312)	(0.7291)
$disc\_index \times \mathbb{1}$[High cig. take-up]	-0.0026	0.0197	0.0153	0.4351
	(0.0251)	(0.3783)	(0.0426)	(1.0203)
Baseline dependent variable	0.147	2.449	0.023	0.592
Observations	402,852	402,852	42,448	42,448

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Note: The unit of analysis is a purchase occasion (beverages in columns (1)–(2), chocolate in columns (3)–(4)), so each sample is conditional on purchasing that category. $disc\_index$ is a store-week, share-weighted measure of discount availability that places greater weight on discounts for more popular products within the category. High cigarette take-up indicates consumers who take cigarette multipack discounts on at least 50% of cigarette trips when such discounts are offered. Standard errors are clustered at the consumer level and reported in parentheses. All specifications include consumer, store, week, and day-of-week fixed effects.

Data Assumptions Calibration

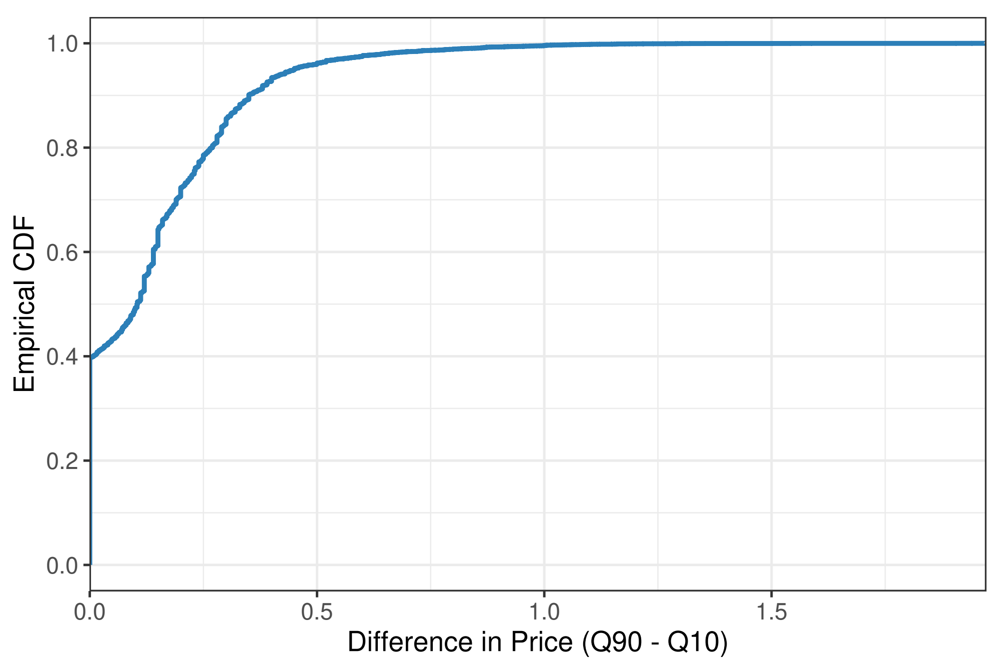

Figure 12: Consumers' Price Dispersion (Q90 - Q10)

Note: This figure shows the empirical distribution of single-pack price dispersion observed among smokers. For each individual, I compute the 10th and 90th percentiles of the single-pack prices across all trips, and take the difference (Q90–Q10). The plot shows the distribution of these differences across smokers.

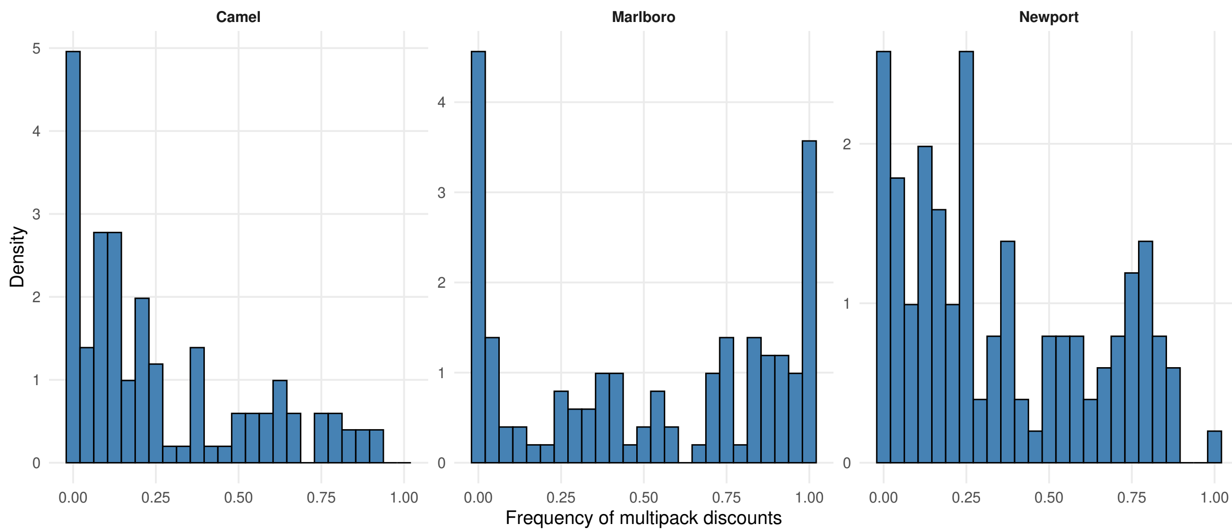

Figure 13: Frequency of Multipack Deals Across Stores

Note: This figure shows the distribution of weekly multipack deals across stores, separately for Marlboro, Camel, and Newport. For each brand-store pair, I compute the share of observed weeks in which a multipack deal was offered. Each panel displays the distribution of these shares across stores for the corresponding brand.

Table 12: Average Days Since Last Trip by Discount Availability

	Average Days Since Last Trip		
	No Discount	**Discount (Total Diff)**	**Discount (Within)**
Value	5.36	-0.15	0.09

Note: This table reports the average number of days since the previous trip, conditional on visiting a store with or without a multipack discount. "Total Diff" reflects the average difference across all individuals, while "Within" captures the average difference within individuals who visit both types of stores.

Exclusion Restriction Check

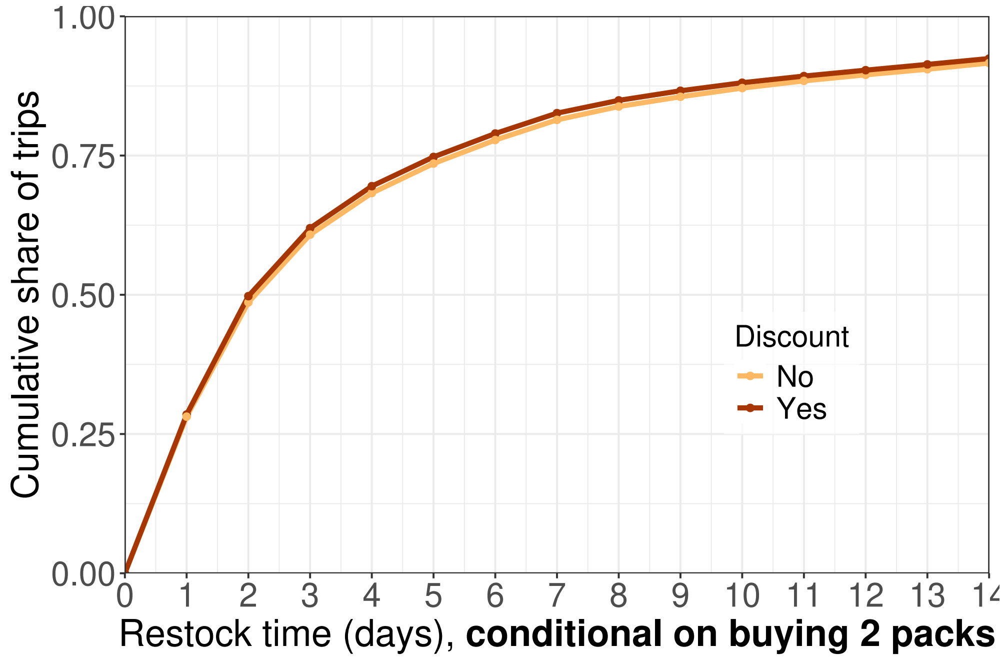

Figure 14: 2-pack Restock Time Distribution-With and Without Discounts

Note: This figure shows the empirical cumulative distribution of restock times (in days) after buying two packs, grouped by whether consumers purchased them on a discount (red curve) or at regular prices (green curve).

Table 13: Effects of Discount on Restock Time, Conditional on Buying Two Packs

	DV: $\log(\tau_{i,t+1})$	
	Coefficient	(Std. Error)
$\mathbb{1}$[discount offered]	0.0042	(0.0061)
Smoker fixed-effects	✓	
Week fixed-effects	✓	
Observations	321,122	

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Note: This table reports results from a regression of $\log(\tau_{i,t+1})$ on $\mathbb{1}$[discount offered]$_{it}$, conditional on observations where $q_{it} = 2$. This analysis tests whether buying two packs on discount results in a different restock time than buying two packs at regular prices. Standard errors are clustered at the individual level and reported in parentheses. All specifications include consumer, store, week, and day-of-week fixed effects.

Table 14: Test for Exclusion Restriction: The Effects of Recent Consumption on Tendency to Visit a Discounted Store

	DV: $\mathbb{1}$[visit discounted store]	
	(1)	(2)
	7-day avg.	14-day avg.
log$(\hat{c})$	0.006***	0.007***
	(0.001)	(0.002)
Baseline dependent variable	0.561	0.561
Observations	835,196	835,196

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.
Note: This table reports results from the regression in Equation 3, which tests whether recent cigarette consumption predicts the likelihood of visiting a store offering multipack discounts. The dependent variable is an indicator for whether the store visited on a given trip offers such a discount. The main covariate of interest is log$(\hat{c})$, where $\hat{c}$ denotes recent consumption. Columns (1) and (2) use 7-day and 14-day moving averages, respectively. Standard errors are clustered at the individual level and reported in parentheses. All specifications include consumer, week, and day-of-week fixed effects.

Discount Take-up by Recent Consumption

Table 15: Discount Take-ups based on Recent Consumption

	DV: $\mathbb{1}[q_{it} = 2]$	
	(1)	(2)
	7-day avg.	14-day avg.
$\mathbb{1}$[discount offered] $\times$ log$(\hat{c})$	0.008***	0.010***
	(0.002)	(0.002)
Baseline dependent variable	0.38	0.38
Consumer fixed-effects	✓	✓
Week fixed-effects	✓	✓
Observations	835,196	835,196

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.
Note: This table reports results from a regression based on the first stage of Equation 2, with the inclusion of an interaction term between discount availability and recent consumption measure. This analysis tests whether recent cigarette consumption moderates discount take-ups. The main covariate of interest is $\mathbb{1}$[discount offered] $\times$ log$(\hat{c})$, where $\hat{c}$ denotes recent consumption. Columns (1) and (2) use 7-day and 14-day moving averages, respectively. Standard errors are clustered at the individual level and reported in parentheses. All specifications include consumer fixed effects, week fixed effects, and day-of-week fixed effects.

Estimating Price Elasticity from the Maryland Tax Increase

On March 14, 2021, Maryland imposed an additional $2 tax on cigarettes. To estimate the price elasticity of demand, I conduct a difference-in-differences analysis at the store–week (s,t) level, using neighboring states as controls (following the approach in Seiler et al. 2021):

$$y_{st} = \alpha_s + \nu_t + \beta \cdot (\text{Maryland}_s \times \text{AfterTax}_t) + \epsilon_{st}$$

where y_{st} is either the log price or log quantity sold, α_s is store fixed effects, and ν_t is week fixed effects.

Table 16: Difference-in-Differences Estimates of Maryland's $2 Tax Increase

	(1)	(2)
	Log price	*Log quantity*
Maryland $\times$ AfterTax	1.864***	$-$31.048***
	(0.019)	(3.107)
Baseline DV	8.00	193.35
% change	23%	-16%
Observations	152,335	152,335

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
Note: This table reports estimates from a difference-in-differences analysis evaluating the impact of Maryland's March 2021 $2 tax increase on cigarette prices and quantities. The estimation compares Maryland stores to those in neighboring states (WV, DE, PA), using store-week level data. All regressions include store and week fixed effects. Standard errors are clustered at the store level and reported in parentheses.