Locked Down and Lighting Up: Using COVID-19 Mobility Restrictions as an Instrument for Cigarette Stock and Smoking Intensity

Daniel Sabey, Don Kenkel, Alan Mathios

June 1, 2026

Abstract

This paper investigates the "temptation effect" – whether increased readily available cigarette stock causally increases smoking intensity. We leverage the unprecedented natural experiment of the COVID-19 pandemic using Nielsen household scanner data, Advan Patterns cell phone data, and USAfacts.org COVID-19 data. An Interrupted Time Series (ITS) model reveals COVID having a sharp and persistent increases in both cigarette stockpiles and smoking rates. Exploration of linked survey data highlights heterogeneous responses related to stress and household characteristics. To isolate the causal effect of stock, we use county-level changes in mobility as an instrumental variable (IV) for changes in cigarette stock. The IV estimates suggest that a one-pack increase in instrumented stock causally increases daily smoking by approximately one cigarette.

Keywords: Smoking, Stockpiling, COVID-19, Instrumental Variables, Mobility Data, Temptation Effect, Nielsen Data

1 Introduction

Cigarette smoking remains a leading cause of preventable death worldwide. While economic factors like price and broader stressors are known correlates of smoking behavior [DeCicca, Kenkel, and Lovenheim, 2022, for a comprehensive and recent review], behavioral economics suggests environmental cues such as product availability may also play a significant role, particularly for addictive goods [Laibson, 2001, Gul and Pesendorfer, 2001]. Does simply having more cigarettes readily available – a larger personal stock – lead individuals to smoke more due to an enhanced "temptation effect"? Answering this question causally is challenging due to the endogeneity of stockpile size. People that smoke more also purchase more. And both choices are endogenous and related to each other directly. This paper proposes to use the COVID-19 pandemic as a natural experiment to address this. The pandemic's onset triggered significant disruptions, including lockdowns and starkly altered shopping patterns, which impacted household spending and economic activity [Baker, Farrokhnia, Meyer, Pagel, and Yannelis, 2020, Goolsbee and Syverson, 2021]. These changes often incentivized consumers to consolidate purchases and increase household inventories, potentially including cigarettes, thereby creating actually plausibly exogenous variation in personal stock levels for study.

We utilize this setting to identify the causal effect of cigarette stock on imputed smoking. First, an Interrupted Time Series (ITS) analysis documents the aggregate impacts on stock and smoking associated with the pandemic onset in 2020. Second, recognizing multiple pandemic pathways, we explore heterogeneity using survey data. Third, to isolate the stock channel, we implement an instrumental variables (IV) strategy using pandemic-induced changes in population mobility (from Advan Patterns data) as an instrument for stock changes. This strategy focuses on the 2019-2020 period capturing the primary shock and relies on controlling for local COVID-19 intensity (from USAfacts.org) and extensive fixed effects to enhance the plausibility of the exclusion restriction.

The findings indicate significant increases in stock and smoking associated with the post-COVID onset period in 2020. Stock increased by nearly 17 cigarettes per household (close to a pack). And smoking increased by approximately one cigarettes per day. The IV analysis indicates a statistically significant causal effect: a one-pack increase in instrumented stock increases imputed daily smoking by approximately one cigarette. This study contributes causal evidence on the stock-consumption link, leveraging unique data and a large-scale shock.

The paper proceeds as follows. Section 2 reviews literature. Section 3 describes data. Section 4 presents the empirical analysis and results. Section 5 discusses findings, and Section 6 concludes.

2 Literature Review

This research sits at the intersection of several literatures.

First is the **economics of addiction** [Becker and Murphy, 1988, Gruber and Köszegi, 2001, Laibson, 2001]. The standard model of rational addiction [Becker and Murphy, 1988] posits that individuals consistently plan to maximize utility over time, even with addictive goods. However, the focus on the "temptation effect" aligns more closely with behavioral models that incorporate time-inconsistent preferences and self-control problems, often building on foundational theories like those presented by Gul and Pesendorfer [2001]. For instance, Gruber and Köszegi [2001] extend addiction models by incorporating hyperbolic discounting, suggesting that individuals may over-consume addictive goods like cigarettes due to present bias, leading to internalities. Work by Laibson [2001] on cue-theories of consumption further explains how environmental cues can trigger impulses and affect self-control, which is highly relevant when considering the immediate availability of a stockpiled addictive good. This theoretical link is strongly supported by empirical work demonstrating how external environmental factors—such as package size, proximity, and visibility—can significantly increase consumption volume, often without the consumer's conscious awareness [Wansink, 2004].

Second, it connects to the literature on **smoking determinants** [Chaloupka and Warner, 2000], adding the role of inventory. Comprehensive reviews such as Chaloupka and Warner [2000] detail the multitude of factors influencing tobacco consumption, including price, income, regulations, and socio-demographic characteristics. The broader economic context of tobacco regulation and behavioral responses is further detailed by DeCicca et al. [2022]. While these determinants are well-studied, the specific impact of consumer-held inventory as a proximate determinant of consumption frequency or quantity remains an area without much exploration simply because the endogeneity of the choice of quantity and smoking.

Third, it relates to studies of **consumer stockpiling** [Chandon and Wansink, 2002, Hendel and Nevo, 2006, Wertenbroch, 1998]. Consumers often stockpile goods due to sales promotions or for convenience, and such inventory can lead to increased consumption rates [Hendel and Nevo, 2006, Chandon and Wansink, 2002]. Chandon and Wansink [2002], for example, found that increased product salience from stockpiles can accelerate consumption. However, this dynamic can be different for "vice" goods. Werten-

broch [1998] demonstrated that consumers may engage in self-control by rationing their purchase quantities of vices, thereby deliberately limiting their home inventory to manage temptation, even if it means forgoing volume discounts. This highlights a crucial tension when an external shock induces stockpiling of such goods.

Finally, it draws on the literature using the **COVID-19 pandemic as a natural experiment** [e.g., Chetty, Friedman, Stepner, and The Opportunity Insights Team, 2020, Lau Andersen, Toft Hansen, Johannesen, and Sheridan, 2020]. The pandemic led to unprecedented shifts in economic activity and consumer behavior, as documented by real-time trackers [Chetty et al., 2020] and analyses of consumer spending responses to shutdowns and stimuli [Lau Andersen et al., 2020]. This period created unique conditions, including widespread stockpiling of essential and non-essential goods [Matsuda, Inakura, and Sakamoto, 2023], which may have forced or encouraged consumers to hold larger-than-usual inventories of products like cigarettes, thereby disrupting typical consumption and self-control patterns.

The most simlar paper to our is a contemporary study by Ho [2025]. It also investigates the causal link between purchase quantity and consumption for cigarettes. Using granular transaction data from convenience stores, Ho [2025] finds that many smokers deliberately forgo multipack discounts as a form of self-control to prevent overconsumption. The paper employs an instrumental variable strategy, using the *availability of a multipack discount* as an instrument for purchasing a larger quantity, and finds that buying an extra pack due to a discount causally increases daily smoking by 52%. This identification strategy, while powerful for testing policy, likely identifies a Local Average Treatment Effect (LATE). The effect is measured *only* for the subset of 'compliers'—smokers who are induced by the discount to purchase two packs. This group likely excludes the most 'sophisticated' consumers who *always* self-ration by buying single packs regardless of price [Ho, 2025]. Our study complements this by analyzing a different, arguably broader treatment. In contrast, our research leverages the COVID-19 pandemic as a natural experiment. This universal, macro-level shock exogenously altered shopping logistics for *all* consumers, plausibly forcing even sophisticated self-rationers to stockpile against their usual preferences. Therefore, the effect we identify may represent a different parameter, arguably closer to an Average Treatment Effect (ATE), as it captures the impact of stockpiling on a population that includes both those who would normally self-ration and those who would not. Interestingly, our IV estimates suggest a smaller causal effect of stock on smoking (approximately a six percent increase) compared to Ho [2025]'s estimate of 50 percent increase in cigarettes per day for compliers. This discrepancy may reflect differences in the populations affected by the two instruments and the nature of the stock increase (forced stockpiling vs. voluntary purchase due to discount).

The **contribution** of this research lies in leveraging this unique shock of the COVID-19 pandemic, with novel data linkages, to provide causal evidence on the stock-consumption relationship for addictive goods like cigarettes. Specifically, it aims to investigate whether and how atypically large personal stockpiles—potentially accumulated contrary to usual self-rationing strategies—affected actual smoking rates through an amplified "temptation effect" stemming from the increased and immediate availability of the addictive good.

3 Data and Sample Construction

This analysis combines detailed household purchase data, granular mobility data, local COVID-19 epidemiological data, and custom survey data to provide a multifaceted view of consumer behavior and its drivers during the pandemic. Each dataset offers unique insights, and their combination allows for a more nuanced understanding of how cigarette purchasing and consumption patterns may have shifted.

3.1 Nielsen Retail Scanner Data and Analysis Sample

The primary data source for observing consumer purchasing behavior is the NielsenIQ Household Panel Data, covering the period from 2015 to 2022. This dataset meticulously tracks household purchases across a wide range of retail outlets, capturing information at the Universal Product Code (UPC) level. For each transaction, it typically includes the date of purchase, products bought, quantity, price paid, and any promotional information, alongside rich demographic information for each participating household. This granularity is crucial for identifying specific cigarette products (brands, pack sizes), calculating expenditure, and inferring purchase frequency and volume, which are key to understanding stockpiling.

To ensure the analysis focuses on a consistent panel of regular cigarette purchasers, thereby minimizing noise from intermittent buyers or households with incomplete reporting, specific inclusion criteria were applied. A "household-year" observation is included if: (1) The household recorded 12 or more distinct cigarette purchase trips within that calendar year, signaling regular engagement in cigarette purchasing; and (2) The household had a span of 250 or more days between their first and last recorded cigarette purchase within that year, indicating consistent participation in the panel and purchasing activity throughout the majority of the year. This filtering process aims to create a robust sample of established cigarette consumers, resulting in 32,340 household-years from 11,388 unique households, as detailed in Table 1. The longitudinal nature of this panel is a significant strength, allowing for the observation of within-household changes in purchasing patterns before and during the COVID-19 pandemic.

Table 1: Yearly Sample Counts (2015-2022)

Year	Unique Households (Freq. from HH-Year)	Household-Day Observations (Freq. from Microdata)
2015	3305	1 171 886
2016	3111	1 105 558
2017	2878	1 018 562
2018	2725	962 897
2019	2733	963 725
2020	2743	966 722
2021	2480	873 100
2022	2158	743 471
Total	22133[a]	7 805 921

[a] Represents the total number of unique household-year observations, not the total number of unique households overall.

3.2 Mobility Data

To capture changes in shopping behavior and physical access to retailers, daily county-level mobility data for the critical 2019-2020 period are sourced from Advan Patterns. Advan Patterns typically provides aggregated and anonymized foot traffic data derived from mobile device locations. The specific measure utilized is the total daily visits to a curated list of points of interest (POIs) whose NAICS codes correspond to likely cigarette retailers. This includes grocery stores, convenience stores, dedicated tobacco stores, and vape shops, among others. This dataset allows for an examination of how public health restrictions,

fear of infection, or other pandemic-related factors influenced the frequency and nature of shopping trips at the county level, which is hypothesized to correlate with shifts towards fewer, larger purchases (i.e., stockpiling).

3.3 COVID-19 and Population Data

Contextual information on the pandemic's local intensity is provided by daily county-level COVID-19 case counts obtained from usafacts.org, a widely used repository for governmental data. County population estimates from the same source allow for the calculation of per capita case rates. To smooth daily fluctuations and capture the prevailing local risk environment, COVID-19 intensity is operationalized as a seven-day moving average of new confirmed COVID-19 cases within a county. This epidemiological data is crucial for linking changes in purchasing, mobility, and potentially consumption to the direct and indirect pressures exerted by the pandemic's local severity.

3.4 Survey Data

Recognizing that scanner and mobility data alone cannot fully elucidate the mechanisms behind behavioral changes, this research incorporates data from a custom-designed survey. This survey aims to collect information on individual-level experiences and attitudes that are hypothesized causal pathways or important control variables. Key constructs measured include perceived stress, financial stress, changes in work-from-home status, shifts in indoor smoking habits (which could be influenced by increased time at home), and details about household composition. This information is vital for understanding how individual circumstances and psychological states during the pandemic might have moderated the relationship between stockpiling and consumption, particularly concerning the "temptation effect."

3.5 Synergies from Data Combination in Understanding COVID-19 Complexities

The novel combination of these diverse datasets provides a powerful lens through which to examine the complex behavioral shifts during the COVID-19 pandemic, particularly concerning cigarette consumption. The NielsenIQ data offers a detailed, longitudinal view of what cigarettes households purchased, when, and in what quantities, forming the backbone for identifying stockpiling and inferring consumption changes. Layering Advan Patterns' mobility data allows for the correlation of these purchasing patterns with changes in real-world shopping behavior, such as reduced trips to retailers, potentially driven by lockdowns or personal avoidance, which are hypothesized to induce stockpiling.

Simultaneously, the USAFacts COVID-19 data provides the dynamic local epidemiological context, enabling an analysis of how the severity of the pandemic in a household's county influenced both their mobility and purchasing decisions. For instance, higher local case rates might correlate with decreased mobility and increased stockpiling.

Crucially, the custom survey data offers the potential to unlock the "why" behind observed changes. While scanner data might show an increase in cigarette inventory followed by increased consumption, the survey data on stress, work-from-home status, or altered indoor smoking norms can help to identify whether the "temptation effect" of a larger stockpile was exacerbated by these pandemic-induced conditions. For example, increased stress combined with a readily available larger stockpile at home (due to WFH) could plausibly lead to higher consumption than either factor alone.

By integrating these datasets, this research can move beyond simple correlations. It can explore how macro-level shocks (the pandemic itself, tracked by USAFacts data) translated into changes in meso-level

behavior (county-wide mobility patterns from Advan) and micro-level household decisions (purchases from NielsenIQ), while accounting for individual-level circumstances and psychological factors (from the survey) that mediate these relationships. This multi-layered approach is essential for disentangling the various influences at play during such an unprecedented period and for providing more robust evidence on the causal pathways linking cigarette stockpiles to consumption via the temptation effect.

3.6 Key Variable Definitions

The following variables are central to the analysis:

- $Smoking_{it}$**:** Imputed cigarettes smoked per day for household i following the most recent purchase (trip t).
  - *Clarification Detail:* This variable aims to measure the daily rate of cigarette consumption for a household. The term "imputed" suggests it's not directly observed but calculated from purchase data. A common method for such imputation involves dividing the number of cigarettes purchased on a given trip (or the inferred stock available after that trip) by the number of days until the household makes its next cigarette purchase. For instance, if a household purchases a carton of 200 cigarettes and their next cigarette purchase is 10 days later, the imputed daily smoking rate for that period would be 20 cigarettes per day. This variable serves as the primary dependent variable for assessing changes in smoking intensity.
  - *Units:* Cigarettes per household per day.
  - *Relevance:* Tracks changes in consumption intensity, which may be influenced by the size of the household's cigarette stock, mobility restrictions, COVID-19 stress, and other factors. It's central to testing the "temptation effect."
- $Stock_{it}$**:** Total cigarettes purchased by household i at their most recent shopping trip t.
  - *Clarification Detail:* This variable quantifies the number of cigarettes acquired by household i during a specific shopping event t. It represents the immediate addition to the household's cigarette inventory from that purchase. For example, if a household buys two packs of 20 cigarettes, $Stock_{it}$ would be 40. This is a key independent variable for assessing how the magnitude of a purchase (and thus the immediate availability of a fresh stock) influences subsequent consumption ($Smoking_{it}$).
  - *Units:* Number of individual cigarettes.
  - *Relevance:* Directly measures the size of a cigarette purchase. The hypothesis is that larger values of $Stock_{it}$ (leading to a larger on-hand inventory) might increase subsequent $Smoking_{it}$ due to greater temptation and availability.
- $Mobility_{ht}$**:** Daily visits per capita to relevant retail points of interest in county h at time t.
  - *Clarification Detail:* This variable, derived from Advan Patterns data, serves as a proxy for overall community mobility and the ease of accessing retail locations where cigarettes are typically sold (e.g., grocery stores, convenience stores, tobacco shops). It is calculated by summing all daily visits to these POIs within county h and then dividing by the county's population (or a segment of it, e.g., adult population) to allow for comparison across counties of different sizes. A decrease in this measure could indicate adherence to lockdown orders, fear of infection reducing shopping trips, or other disruptions to normal shopping routines.

  - *Units:* Daily visits per person (or often scaled, e.g., visits per 1,000 people) in county h on day t.
  - *Relevance:* Captures changes in shopping behavior at a geographic level that might influence purchasing strategies (e.g., fewer trips leading to larger basket sizes and stockpiling). It's a crucial indicator of how pandemic-related restrictions and behaviors impacted access to and interaction with the retail environment.

- $CovidMeasure_{ht}$**:** Average new COVID-19 cases per 1000 people (over the past 7 days) in county h at time t.
  - *Clarification Detail:* This variable quantifies the recent local intensity of the COVID-19 pandemic. It's constructed by taking the number of new COVID-19 cases reported in county h over the 7 days leading up to and including day t, averaging this to a daily figure, and then normalizing it by the county population (expressed per 1,000 people). The 7-day averaging helps to smooth out short-term fluctuations in case reporting.
  - *Units:* Average daily new COVID-19 cases per 1,000 population.
  - *Relevance:* Represents the prevailing local health risk and pandemic severity. Higher values might correlate with increased stress, stricter adherence to public health guidance, reduced mobility, and potentially altered consumption patterns for goods like cigarettes. It allows for examining how the direct impact of the pandemic's intensity in the local environment relates to smoking and purchasing behaviors.

This detailed operationalization of key variables will enable a nuanced exploration of the interplay between cigarette purchasing, stockpiling, local pandemic conditions, mobility, and imputed daily consumption.

4 Empirical Analysis and Results

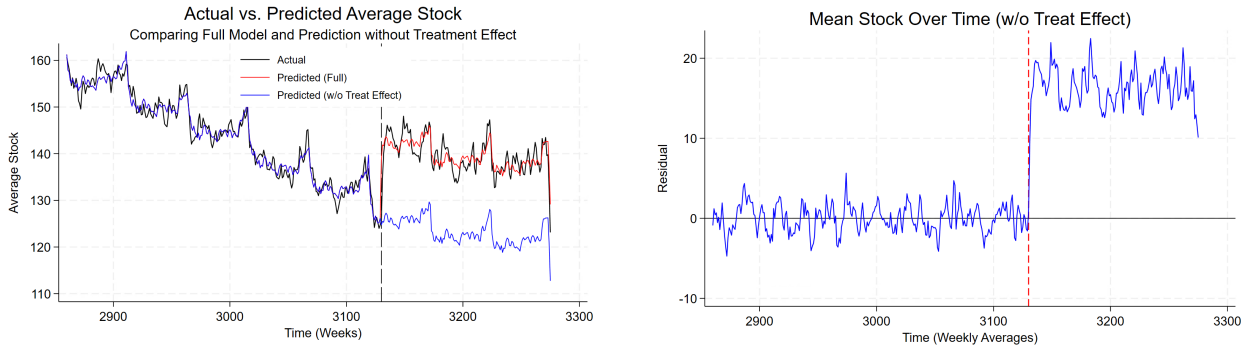
4.1 Overall Impact of COVID-19 (ITS Analysis)

To assess the average impact associated with the pandemic onset period in 2020, an Interrupted Time Series (ITS) model is employed. This model controls for underlying seasonality and time-invariant household characteristics through fixed effects, allowing for a clearer estimation of the pandemic's distinct effect. Let Y_{iht} be the outcome of interest (either $Stock_{it}$ or $Smoking_{it}$). The model is specified as:

$$Y_{iht} = \beta_0 + \beta_1 \mathbb{I}(DOY_t \geq 72) + \beta_2[\mathbb{I}(Year_t = 2020) \times \mathbb{I}(DOY_t \geq 72)] + \sum_{y=2015}^{2022} \delta_y Trend_{yt} + \alpha_w + \alpha_y + \alpha_i + \epsilon_{iht}$$

Where DOY_t is the day of the year, and $\mathbb{I}(DOY_t \geq 72)$ is an indicator function that equals one for observations on or after the 72nd day of the year (approximately March 12th, corresponding to the period when widespread awareness and initial major responses to COVID-19, such as the U.S. national emergency declaration, began). The term β_1 captures any average shift in the outcome variable that occurs after DOY 72 across all years in the sample, accounting for typical seasonal changes around this time of year. The key parameter of interest is β_2, which estimates the differential level change in the outcome specifically for the post-onset period (DOY $\geq$ 72) in the year 2020, compared to the same period in other years. $Trend_{yt}$ represents year-specific linear time trends, included to capture any distinct secular changes in purchasing or smoking behavior unique to each year, beyond the acute seasonal or event-driven shifts.

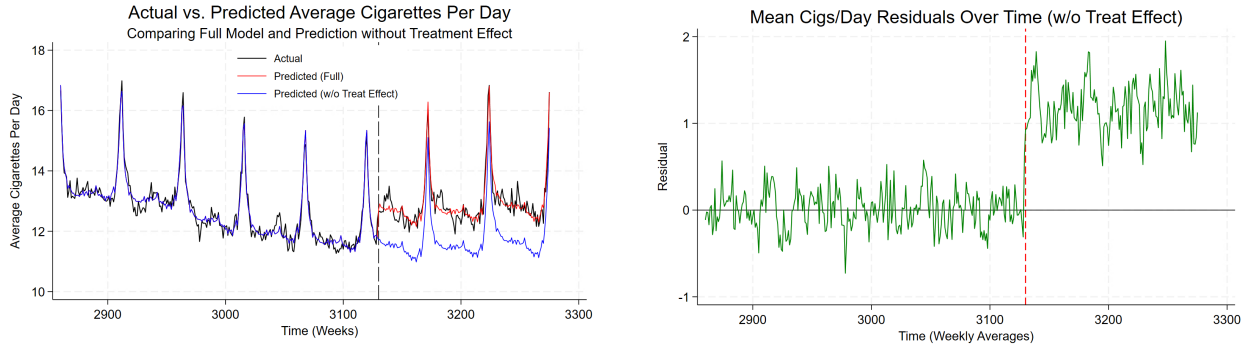
Fixed effects are included for week-of-the-year (α_m, to control for fine-grained seasonality), year (α_y, to capture year-specific baseline levels), and household (α_i, to control for all time-invariant household characteristics). Standard errors are clustered at the household level to account for potential intra-household correlation in the errors over time.



(a) Actual vs. Predicted (b) Residuals

Figure 1: ITS Model Fit and Residuals for Average Stock Purchased Per Trip Over Time

Figure 1 visually presents the ITS model's performance for stock purchased per trip. Subfigure 1a (Actual vs. Predicted) demonstrates that the model tracks the underlying seasonal patterns and pre-existing trends in stock purchases reasonably well prior to the March 2020 discontinuity. This fit lends confidence to the model's ability to project the counterfactual scenario (i.e., what would have been expected in 2020 absent the pandemic-specific shock). The clear upward deviation of actual purchases from predictions around March 2020 highlights the pandemic's impact. Subfigure 1b (Residuals) shows the errors of the model over time without including the treatment effect. In other words, this is visually showing what the counterfactual would have looked like if we only used the pre-period to predict the counterfactual. The persistence of the increase in stock is notable. It still seems to be significantly higher at the end of 2023.



(a) Actual vs. Predicted (b) Residuals

Figure 2: ITS Model Fit and Residuals for Imputed Daily Smoking Over Time

Figure 2 shows a similar visual analysis for imputed daily smoking, again illustrating a discernible increase in actual smoking compared to model predictions around the pandemic's onset. Here too we see the same sticky changes in smoking behaviors through the end of 2023. Table 2 quantifies these average changes. Associated with the post-onset period in 2020 (the β_2 coefficient), there is an average

differential increase of 1.177 imputed cigarettes smoked per day ($p<0.01$) and a 16.385 cigarette increase in stock purchased per trip ($p<0.01$), relative to comparable periods in other years and conditional on the extensive set of controls. To put this in perspective, if the pre-pandemic average stock purchase was 140 cigarettes, a 16.4 cigarette increase represents a substantial percentage rise. Similarly, an increase of over one cigarette per day per household is a meaningful shift in consumption intensity.

Table 2: Interrupted Time Series (ITS) Results: Impact of Post-COVID 2020 Period

	(1)	(2)
OUTCOME VARIABLES	Cigs Per Day	Stock
Post-Covid 2020 Interaction (β_2)	1.177***	16.385***
	(0.203)	(2.134)
Observations (Household-Days)	[7,805,921]	[7,805,921]
Fixed Effects	HH, Week, Year	HH, Week, Year
Year-Specific Trends	Yes	Yes
Post-Onset Main Effect	Yes	Yes

Notes: Coefficients shown are for the interaction term $[\mathbb{I}(Year_t \geq 2020) \times \mathbb{I}(DOY_t \geq 72)]$ from the ITS model. Standard errors in parentheses, clustered at the household level. *** $p<0.01$. Model includes specified fixed effects and trends.

Interestingly, visual inspection of raw data trends suggests that the COVID impact on purchasing and smoking levels might be large and that these changes appear sticky, with levels remaining elevated even as the initial severity of the pandemic waned. In fact this was true through the end of the data we have through the end of 2023. This observation motivates further research into the longer-term dynamics of these behavioral shifts. It intuitively makes sense that smokers were exogenously pushed to smoke more and became addicted to the new level of consumption so that level of consumption start for 3 years after the initial shock. This pattern also implies an increase in the cost on health of policies that unintentionally increase smoking.

Naively, using the average treatment effects from the ITS, one could calculate a Wald-style estimate of the relationship between the change in smoking and the change in stock ($1.177/16.385 \approx 0.0718$). This would imply that, on average, for each additional cigarette added to a purchase trip during this shock period, daily smoking increased by approximately 0.07 cigarettes, or that a one-pack (20 cigarettes) increase in stock purchased per trip was associated with an average increase of about 1.44 cigarettes smoked per day (20×0.0718). This estimate is considered "naive" as it relies on aggregate changes and does not account for individual-level heterogeneity or confounding variables that could simultaneously drive both stock and smoking (e.g., stress, changes in work patterns). It also assumes the average effect of the pandemic on stock and its average effect on smoking are linked in a way that perfectly reflects the marginal effect of stock on smoking, an assumption our subsequent IV strategy aims to address more directly.

4.2 Heterogeneity and Alternative COVID-19 Pathways

The COVID-19 pandemic was a multifaceted shock, likely influencing behavior through various channels beyond just altered shopping logistics leading to larger stockpiles. Factors such as increased stress, financial uncertainty, significant changes in daily routines (e.g., work-from-home mandates), and shifts in household dynamics could all independently or interactively affect smoking behavior. Examining heterogeneity in the pandemic's impact across different subgroups, based on their contemporaneous survey responses, can provide valuable insights into these alternative pathways and the varied experiences of households. For this analysis, the same Interrupted Time Series (ITS) model specified earlier is estimated separately for subgroups defined by their survey responses to questions about stress, financial stress, indoor smoking habits, and the presence of school-aged children. Respondents were categorized based on their self-reported levels for these factors (e.g., for stress, into "No Stress," "Some Stress," and "High Stress" based on their answers to relevant survey items).

Table 3: Heterogeneity in COVID-19 Effects Across Survey Subgroups

	Dependent Variable (Differential Change Post-COVID-19)	
	Cigs Per Day	Stock
Panel A: Stress Groups		
No Stress	2.614***	17.853*
	(0.864)	(9.142)
Obs (HH)	[272,723] (261)	[272,723] (261)
Some Stress	0.425	8.363
	(0.440)	(5.407)
Obs (HH)	[824,078] (848)	[824,078] (848)
High Stress	0.661	15.281**
	(0.770)	(6.584)
Obs (HH)	[378,029] (437)	[378,029] (437)
Panel B: Financial Stress Groups		
No Financial Stress	1.854**	17.605**
	(0.779)	(949)
Obs (HH)	[555,872] (545)	[555,872] (545)
Some Financial Stress	-0.037	4.282
	(0.440)	(4.261)
Obs (HH)	[616,773] (655)	[616,773] (655)
High Financial Stress	1.056	17.363**
	(0.895)	(7.773)
Obs (HH)	[302,185] (346)	[302,185] (346)
Panel C: Indoor Smoking Groups		
No Inside Smoking	1.052*	13.393**
	(0.574)	(5.550)
Obs (HH)	[767,594] (845)	[767,594] (845)
Some Smoking Inside	-0.702	9.020
	(0.935)	(9.603)
Obs (HH)	[294,413] (303)	[294,413] (303)
Smoking Inside	1.653***	10.861*
	(0.628)	(5.961)
Obs (HH)	[412,823] (398)	[412,823] (398)
Panel D: School-Aged Children Groups		
School Aged Kids	2.178***	18.129***
	(0.684)	(6.098)
Obs (HH)	[207,537] (242)	[207,537] (242)
No School Aged Kids	0.668	10.833***
	(0.444)	(4.357)
Obs (HH)	[1,267,293] (1,304)	[1,267,293] (1,304)

Notes: Coefficients estimate the differential change post-COVID (specifically, $[\mathbb{I}(Year_t = 2020) \times \mathbb{I}(DOY_t \geq 72)]$) for the specified subgroup, based on the ITS model. Standard errors in parentheses, clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Table 3 reveals some variation in the estimated impact of the post-COVID 2020 period (β_2) across these subgroups. However, while the point estimates differ, formal statistical tests indicate that most of these differences between subgroups are not statistically significant at conventional levels. Nonetheless, the patterns that emerge, even if not definitively distinct, are suggestive of different causal pathways through which the pandemic may have influenced household behavior.

Several patterns are noteworthy:

- **Stress and Financial Stress:** The point estimates for households reporting *No Stress* or *No Financial Stress* are among the largest and most statistically significant for increases in both imputed daily smoking and stock purchased per trip. For instance, the "No Stress" group's estimated increase of 2.614 cigarettes per day is statistically different from the smaller, insignificant effect for the "Some Stress" group ($p = 0.034$). This might seem counterintuitive, but it could suggest that these households, perhaps less burdened by psychological or economic precarity, had a greater capacity to alter consumption patterns (e.g., more disposable income for stockpiling). Conversely, the lack of a strong statistical difference between the "High Stress" and "No Stress" groups suggests a complex relationship that is not a simple linear function of stress.

- **High Stress Groups:** Interestingly, the "High Stress" (Stock: 15.281) and "High Financial Stress" (Stock: 17.363) groups also show notable point estimates for increases in stock purchased per trip. While the corresponding estimates for daily smoking are not statistically significant on their own, they align with theories that stress can be a trigger for increased substance use. The lack of statistical power to distinguish these groups from others highlights the difficulty in isolating this specific mechanism.

- **Indoor Smoking:** The behavioral response appears to differ based on pre-existing household smoking norms. Households that reported "Smoking Inside" pre-pandemic show a significant increase in imputed daily smoking (1.653 cigarettes/day, $p < 0.01$). This effect is statistically different from the negative point estimate for the "Some Smoking Inside" group ($p = 0.033$). This could indicate that the increased time spent at home, combined with an existing norm of indoor smoking, amplified the impact of any stockpiled cigarettes.

- **School-Aged Children:** Households with "School-Aged Kids" demonstrate a pronounced increase in both smoking (2.178 cigarettes/day, $p < 0.01$) and stock per trip (18.129 cigarettes, $p < 0.01$). The difference in imputed daily smoking is marginally significant ($p = 0.077$) when compared to households without school-aged children. This substantial effect could reflect a combination of factors: significantly altered daily routines due to school closures, increased parental stress not fully captured by the general stress variable, and more time spent in the home where stockpiles are accessible.

These varied patterns, while generally not statistically distinct from one another, point toward the complex and heterogeneous nature of the pandemic's impact. The same broad shock (COVID-19 onset) appears to have been filtered through different household experiences, leading to diverse behavioral responses. This observed heterogeneity, combined with the statistical ambiguity in the differences between groups, underscores the challenge of disentangling causal effects using simple subgroup analysis. It therefore strongly motivates the need for an instrumental variables approach to isolate the specific causal effect of changes in cigarette stock on consumption, attempting to control for these concurrent influences like stress or routine disruption.

4.3 Isolating the Stock Effect (Instrumental Variables Analysis)

To disentangle the causal effect of changes in cigarette stock ($Stock_{it}$) from other confounding pandemic-related factors (e.g., stress, changes in daily routines) that might simultaneously affect both stockpiling behavior and smoking intensity ($Smoking_{it}$), we employ a two-stage least squares (2SLS) instrumental variables (IV) strategy. County-level mobility per capita ($Mobility_{ht}$) is used as an instrument for household cigarette stock purchases. The underlying logic is that broad, county-level shifts in public mobility—driven by exogenous factors like official lockdown orders, business operating restrictions, or widespread public health advisories during the pandemic—directly constrained households' typical shopping routines and frequency. Such constraints would incentivize households to make fewer but larger shopping trips to minimize exposure or comply with restrictions, leading to an increase in the quantity of cigarettes purchased per trip ($Stock_{it}$). These mobility shifts are arguably exogenous to an individual household's idiosyncratic changes in smoking preferences or stress levels, once common factors are controlled for.

1. First Stage Regression:

$$Stock_{it} = \alpha_0 + \alpha_1 Mobility_{ht} + \pi CovidMeasure_{ht} + \alpha_m + \alpha_i + \alpha_h + u_{iht}$$

2. Second Stage Regression:

$$Smoking_{it} = \gamma_0 + \gamma_1 \widehat{Stock}_{it} + \theta CovidMeasure_{ht} + \alpha_m + \alpha_i + \alpha_h + v_{iht}$$

Where $\widehat{Stock}_{it}$ is the predicted value of $Stock_{it}$ from the first stage. Fixed effects are included for month-of-year (α_m, to capture seasonality in purchasing and smoking), household (α_i, to control for all time-invariant household characteristics like baseline preferences or income), and county (α_h, to account for time-invariant county-level factors such as baseline retail environment or local smoking norms). Both stages also control for the contemporaneous local COVID-19 intensity ($CovidMeasure_{ht}$).

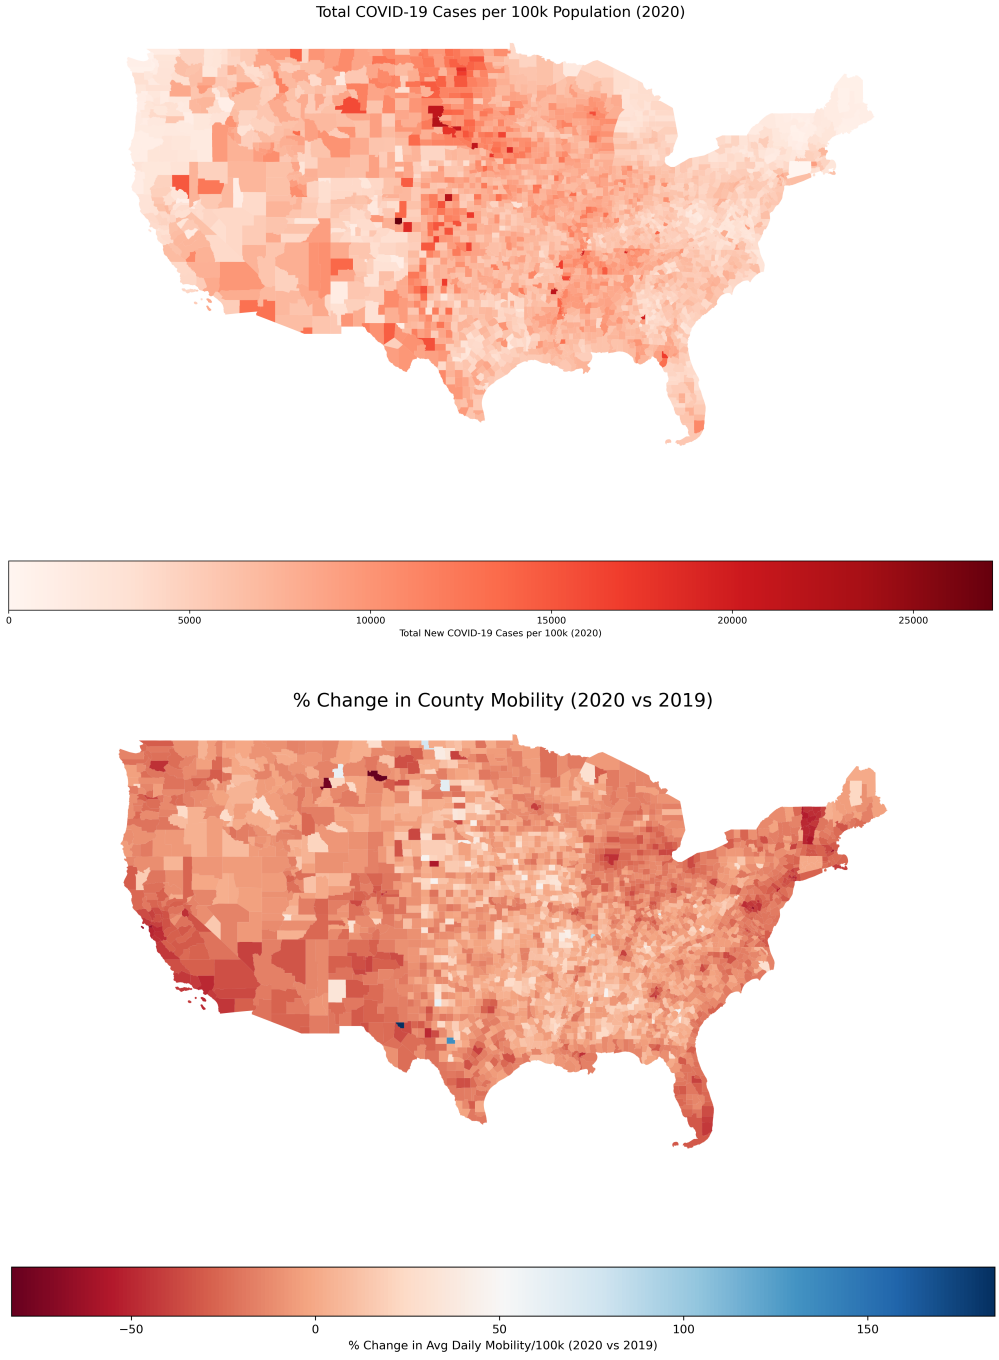


Figure 3: Maps illustrating COVID-19 severity and mobility changes.

Figure 3 illustrates the geographical variation across U.S. counties in both the instrument ($Mobility_{ht}$, represented by changes in mobility) and a key control variable, COVID-19 case intensity ($CovidMeasure_{ht}$). It also shows substantial regional differences in how mobility patterns changed from 2019 to 2020, likely reflecting varied local policy responses (e.g., stringency of lockdowns) and potentially different levels of public adherence. We also see the variation in COVID-19 case intensity. Visually, there isn't a perfect one-to-one correlation between high COVID intensity and the largest drops in mobility; some areas with high case rates still maintained relatively higher mobility than others with lower case rates but perhaps stricter local ordinances or different population responses. This imperfect correlation is advantageous for the IV strategy: it suggests that $Mobility_{ht}$, while expectedly related to the pandemic, captures distinct sources of variation beyond just the raw local COVID case intensity. This distinct variation, plausibly linked more to the logistical challenges and altered costs of shopping, is what the instrument leverages after directly controlling for $CovidMeasure_{ht}$ in the regressions. The substantial county-to-county variation in these plausibly policy-induced mobility changes is key to identifying the effect of stock on consumption, as these changes are arguably exogenous to individual smokers' prior choices and evolving smoking preferences during the pandemic.

The validity of the IV strategy hinges on the exclusion restriction: $Mobility_{ht}$ must affect $Smoking_{it}$ only through its effect on $Stock_{it}$, conditional on the included controls and fixed effects. The inclusion of Household (α_i), County (α_h), and Month-of-Year (α_m) fixed effects is crucial, as it controls for a vast array of unobserved time-invariant factors at these levels (e.g., stable household preferences, baseline local retail density, ingrained seasonality in smoking). Furthermore, explicitly controlling for local COVID intensity ($CovidMeasure_{ht}$) is vital, as it ensures that the mobility instrument isn't merely reflecting direct pandemic impacts on smoking that operate through channels like heightened stress or health fears driven by local case rates. By conditioning on these rich controls, the remaining variation in mobility is more plausibly linked to logistical constraints or opportunities affecting purchasing behavior (and thus the quantity of cigarettes bought per trip, $Stock_{it}$) rather than directly influencing smoking through other unobserved, dynamic, within-county channels. For the exclusion restriction to be violated, an unobserved factor would need to vary dynamically within households and counties over months, directly affect smoking decisions, be correlated with county-level mobility changes, and yet be uncorrelated with our direct measure of local COVID intensity – a relatively stringent set of conditions. While such a confounder cannot be definitively ruled out, the extensive set of controls makes the exclusion restriction considerably more plausible.

To provide further reassurance against the primary confounding concerns, we created different measures of the instrument's components: the average mobility before COVID, the average mobility after COVID, and the difference between the two in each county. We then regress each of the versions of our instrument on key household characteristics that could drive unobserved changes in smoking with the subsample of people that we have a survey response for. The regression results presented in Table 4 show that across the three instrument components and four potential confounders, 11 out of the 12 regression coefficients are not statistically different from zero at the conventional p¡0.05 level. This is particularly reassuring because the main hypothesized confounding path that bypasses the stock mechanism involves psychological or financial stress directly affecting smoking. Since the instrument is independent of these direct measures of stress, the overall threat to the exclusion restriction is significantly diminished.

Although the mobility difference measure is statistically correlated with having school-age children ($\beta = 115.436, p = 0.036$), we do not think it is not a critical threat, for two main reasons. First, the presence of school-aged children is largely a time-invariant household characteristic, and its base effect on smoking is absorbed by the Household Fixed Effect (α_i) in the IV second stage. Second, the observed correlation would be a problem if it causes more smoking independently, however the main pathway we

Table 4: Regression Results: Instrument Components and Correlation with Potential Confounders

	Dependent Variable: County Mobility (Visits per 100k)		
Potential Confounder	**(1)** **Pre-COVID**	**(2)** **Post-COVID**	**(3)** **Change**
CovidStress	-86.992	-135.104	48.606
	(82.123)	(82.677)	(40.001)
	$[p = 0.290]$	$[p = 0.102]$	$[p = 0.225]$
CovidFinStress	-106.003	-82.995	-34.417
	(68.988)	(67.056)	(36.077)
	$[p = 0.125]$	$[p = 0.216]$	$[p = 0.340]$
HomeSmokingPolicy	-1.726	33.894	8.779
	(56.387)	(51.704)	(27.740)
	$[p = 0.976]$	$[p = 0.512]$	$[p = 0.752]$
SchoolAgeKids	169.514	18.871	115.436^{**}
	(125.188)	(129.652)	(54.991)
	$[p = 0.176]$	$[p = 0.884]$	$[p = 0.036]$
Observations	1,540	1,195	1,189

Notes: Regression results showing the correlation between the components of the mobility instrument and various potential confounders. The main coefficient is listed first, followed by the robust standard error in parentheses, and the p-value in square brackets. Column (1) reports results using pre-COVID mobility as the dependent variable. Column (2) uses post-COVID mobility. Column (3) uses the difference in mobility ($\text{Mobility}_{post} - \text{Mobility}_{pre}$). Each column represents a separate set of regressions. Mobility values are in units per 100,000 people.
Significance: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

would expect would again be through the stress pathway; School aged kids resulted in more stress and thus more smoking. However, the result is not correlated with stress, so we do not think this should be interpreted as a possible red flag. Instead we think is best interpreted as evidence of a stronger-than-average first stage. The correlation suggests that families with children tended to live in counties that experienced the largest negative changes in mobility. This likely reflects that households with school-aged children faced a disproportionately higher logistical burden during lockdowns (due to remote schooling and supervision), which intensified the need to consolidate shopping. The correlation thus indicates that the instrument is simply more relevant and powerful for this specific subgroup, thereby reinforcing the plausibility of the causal stock effect we seek to isolate rather than violating the exclusion restriction. These results are thus reassuring that the instrument is plausibly uncorrelated with major potential time-varying drivers of smoking behavior not operating through the intended stock channel.

Table 5: IV First Stage Results

Dependent Variable:	Stock ($Stock_{it}$)
Mobility ($visits_per_person_{ht}$)	[-285.599] (58.718)
Observations	1,021,071
R-squared (Centered)	0.0691
Kleibergen-Paap rk LM Stat (Underid.)	27.795
P-value for Underid. Test	0.0000
Kleibergen-Paap rk Wald F Stat (Weak id.)	23.638
Stock-Yogo Crit Val (10% Max Rel Bias)	16.38
Fixed Effects	HH, County, Month

Notes: Results from First Stage IV regression. Standard errors clustered by household in parentheses. Model includes Household, County, and Month fixed effects and controls for $CovidMeasure_{ht}$. The K-P Wald F-statistic significantly exceeds conventional thresholds, indicating strong instrument relevance.

Table 5 presents the first stage results for the IV estimation using the 2019-2020 sample. The coefficient on $Mobility_{ht}$ (visits per person) is -285.599 and is highly statistically significant. The negative sign indicates that decreased county-level mobility is strongly associated with an increase in the number of cigarettes households purchase per trip ($Stock_{it}$). This aligns with the hypothesized mechanism: as general mobility reduces, making shopping trips potentially more costly or risky, households tend to consolidate purchases, buying larger quantities per trip. The Kleibergen-Paap rk Wald F-statistic for weak identification is 23.638. This value strongly rejects the null hypothesis of a weak instrument and exceeds the Stock-Yogo critical value for 10% maximal relative bias (16.38), indicating that county-level mobility is a strong and relevant predictor of household cigarette stock purchases in this sample, satisfying a key requirement for a valid instrument.

Table 6: Comparing OLS and IV Estimates of Stock Effect on Smoking Per Day

Dependent Variable:	Cigarettes Smoked Per Day ($Smoking_{it}$)				
VARIABLES	(1)	(2)	(3)	(4)	(5)
Specification:	OLS	OLS	2SLS (IV)	2SLS (IV)	2SLS (IV)
Sample Period:	2019-2020	2019-2020	2019-2020	2019-2020	2019 Only
Stock ($\widehat{Stock}_{it}$ in IV)	0.063***	0.068***	0.046***	0.040***	-0.036
	(0.005)	(0.005)	(0.013)	(0.012)	(0.171)
Observations	1,921,071	1,921,071	1,921,071	1,921,071	954,813
Controls:					
COVID Cases Control	No	Yes	Yes	No	No
Household FE	No	Yes	Yes	Yes	Yes
County FE	No	Yes	Yes	Yes	Yes
Month FE	No	Yes	Yes	Yes	Yes

Notes: Robust standard errors clustered by household in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Columns (1) and (3) represent the main OLS and IV estimates on the 2019-2020 sample. Column (4) shows the IV estimate without controlling for COVID cases per 1000. Columns (2) and (5) show OLS and IV estimates on the 2019 pre-period sample only. The instrument for IV models is county-level visits per person ($Mobility_{ht}$).

Table 6 presents the main causal estimates from the second stage, comparing OLS and IV results. The OLS estimate using the 2019-2020 sample without household, county, month, and year fixed effects, and without controlling for COVID cases, suggests that a one-cigarette increase in stock purchased is associated with a 0.063 increase in imputed daily smoking. This is the totally naive approach, or in other words is pure correlation. Interestingly, as we control for COVID, and month and household, we see a slight increase in the estimate. However, clearly both estimates are biased. Namely because people who smoke more also purchase more. But it gives us a baseline to compare to. Of course, it may also be biased if unobserved factors (e.g., increased stress during the pandemic) simultaneously lead individuals to both buy more cigarettes and smoke more. Our preferred 2SLS IV estimate (Column 3), using $Mobility_{ht}$ as an instrument and including the full set of fixed effects and the COVID cases control, is 0.046 and remains highly statistically significant (p<0.01). This implies that a one-unit increase in instrumented cigarette stock leads to a causal increase of approximately 0.046 imputed cigarettes smoked per day. Equivalently, obtaining an extra pack (20 cigarettes) via mobility-induced stock changes causally increases daily smoking by nearly one full cigarette ($20 \times 0.046 = 0.92$). And given the COVID-19 context, that means the 16.385 cigarette increase in stockpiles was causally connected to a 0.753 cigarettes smoked per day ($16.385 \times 0.046 = 0.753$). That is 64 percent of the total increase in smoking, leaving the other 36 percent to be connected to other causes like COVID-19 related stress.

The IV estimate (0.046) is smaller than the OLS estimate (0.063) and the naive Wald-style estimate from the ITS results (0.0718). This suggests that the OLS and Wald estimates were likely biased upwards, consistent with the presence of positive omitted variable bias. For example, heightened anxiety during the pandemic might have independently driven individuals to both purchase larger stocks of cigarettes and also increase their smoking frequency; OLS would conflate these effects, while the IV attempts to isolate

the causal impact of the stock itself.

Column (4) presents the IV estimate for 2019-2020 without directly controlling for $CovidMeasure_{ht}$. The coefficient on stock (0.040) is slightly smaller but remains statistically significant. The similarity of this estimate to that in Column (3) suggests that the instrumented stock effect is robust to the inclusion of the direct COVID measure in the second stage, though its inclusion is theoretically preferred to account for any direct pathway from local pandemic severity to smoking that is not mediated by stock.

Column (5) present and IV estimate using data from the 2019 pre-pandemic period only. The IV estimate for 2019 is -0.036 and is not statistically significant (given the large standard error of 0.171). This "placebo" test is important: it suggests that the instrument (county-level mobility) does not have a significant relationship with smoking through stock in a period without the pandemic-induced exogenous mobility shocks. This strengthens the argument that the significant IV results for 2019-2020 are indeed capturing effects driven by the unique conditions of the pandemic, where mobility changes acted as a strong shifter of purchasing behavior.

These results provide strong evidence supporting the "temptation effect": larger, exogenously-driven personal stockpiles of cigarettes appear to causally increase smoking intensity.

5 Discussion

This paper investigates the causal effect of cigarette stock on consumption, a relationship often confounded by endogeneity, by leveraging the unique circumstances of the COVID-19 pandemic as a natural experiment. The descriptive Interrupted Time Series (ITS) analysis documents statistically significant average increases in both cigarette stockpiling (approximately 16 more cigarettes per purchase trip) and imputed daily smoking (approximately 1 more cigarette per day per household) associated with the post-onset period of the pandemic in 2020, relative to previous years and controlling for various trends and fixed effects. This change lasts until at least 2023 when our data ends. COVID has impacted and my still be impacting smokers to purchase more cigarettes and to smoke more.

Further heterogeneity analysis using survey data reveals that these average impacts varied considerably across households, with factors like reported stress levels, financial stress, indoor smoking habits, and the presence of school-aged children associated with different magnitudes of change in purchasing and smoking. This highlights the complex alternative pathways through which the pandemic influenced behavior and underscores the need to isolate the specific effect of stock.

The core of the paper's contribution lies in the Instrumental Variables (IV) strategy. Using county-level changes in public mobility (from Advan Patterns) during 2019-2020 as an instrument for changes in household cigarette stock purchases, while controlling for local COVID-19 intensity (from USAfacts.org) and an extensive set of household, county, and month-of-year fixed effects, yields a strong first-stage relationship (F-statistic = 23.6). The resulting second-stage IV estimate indicates a statistically significant and positive causal effect of stock on imputed daily smoking: a one-pack (20 cigarettes) increase in instrumented stock is estimated to increase daily smoking by approximately 0.92 cigarettes.

The IV estimate (0.046 per cigarette in stock) is somewhat smaller than the comparable OLS estimate (0.063) and the naive Wald estimate from the ITS results (0.0718). This suggests that OLS and simpler correlational approaches may indeed overestimate the true causal effect of stock on smoking, likely due to unobserved confounders (e.g., a general increase in anxiety during the pandemic might lead individuals to both buy more cigarettes and independently smoke more). The IV approach, by isolating variation in stock purchases driven by plausibly exogenous shifts in shopping mobility, appears to mitigate this upward bias.

The plausibility of the exclusion restriction—that county-level mobility only affects household smoking through its impact on household cigarette stock, conditional on controls—is central to the IV strategy's validity. This assumption is strengthened by the rich set of controls employed. Household and county fixed effects account for a wide range of unobserved time-invariant characteristics at these levels. Crucially, controlling for local COVID-19 case intensity helps to ensure that the mobility instrument is not merely proxying for the direct psychological or behavioral effects of local pandemic severity (such as fear-induced self-isolation that might also independently affect stress and smoking). For a violation to occur, a remaining unobserved confounder would need to be dynamic within households and counties, directly influence smoking patterns, be correlated with the county-level mobility shifts, and yet be uncorrelated with our measure of local COVID intensity—a relatively demanding set of conditions. While no instrument is perfect, the strategy employed here offers a credible approach to identifying the causal parameter of interest.

The finding that larger readily available stocks causally increase cigarette consumption has several potential policy implications:

- **Purchase Quantity Limits and Availability:** While direct restrictions on purchase quantities (e.g., limiting sales to single packs) could be controversial and might lead to unintended consequences like increased shopping trips or illicit trade, the findings highlight a behavioral mechanism that policymakers should be aware of. During public health crises or other situations anticipated to trigger widespread stockpiling, public health campaigns could include messaging that cautions smokers about the potential for increased consumption due to larger home inventories and suggests strategies for managing this risk.

- **Cessation Strategies:** For individuals attempting to quit or reduce smoking, this research underscores the behavioral principle of limiting the availability and accessibility of the addictive substance. Smoking cessation programs and healthcare providers could more actively incorporate advice on conscious purchase rationing (e.g., buying only single packs, avoiding "stocking up" even if there are sales) as a practical self-control strategy to complement existing pharmacological and behavioral therapies. This aligns with Wertenbroch's (1998) findings on self-rationing.

- **Taxation Policy and Stockpiling:** Anticipation of significant cigarette tax increases often leads to pre-tax stockpiling. The results suggest that the temptation effect from these accumulated stockpiles could temporarily offset some of the intended consumption-reducing impact of the tax hike itself. This doesn't negate the overall effectiveness of taxes but suggests that the dynamics around their implementation might need to consider this behavioral response.

- **Retail Environment:** While not directly studied, if ease of purchase (related to retailer density, a factor controlled for by county FEs in part) affects stock levels, then policies related to the tobacco retail environment might have indirect effects on consumption through this stock-temptation channel.

Several limitations should be acknowledged:

The primary outcome, daily cigarettes smoked, is imputed from purchase data rather than directly observed. This method, while standard in the absence of direct consumption tracking, relies on assumptions such as consistent consumption rates between purchases and that all household cigarettes are sourced from scanned purchases. If households obtain cigarettes from unrecorded sources (e.g., gifts, illicit markets) or if their daily consumption is highly erratic in ways not tied to purchase cycles, the imputed measure will contain error. This could potentially bias the estimated effect, although the direction is not immediately obvious.

The pandemic was an extraordinary global event characterized by a unique confluence of widespread health fears, significant economic disruption, unprecedented work-from-home arrangements, and heightened psychological stress. While this provided a valuable natural experiment for identifying the stock-consumption link, the magnitude of the "temptation effect" observed might be amplified by or specific to this high-stress, altered-routine environment. Caution is warranted when generalizing the quantitative findings to stockpiling behaviors induced by more typical circumstances, such as responses to price promotions or routine bulk shopping, where the same level of concurrent stressors might not be present. That being said, we do think that the unique natural experience that is the COVID-19 pandemic gives us a rare chance to get at an exogenous change in stockpiling.

County-level mobility is an aggregate measure. While it serves as a strong instrument for average changes in purchasing logistics within a county, individual households' actual mobility and shopping constraints might vary from the county average. Household fixed effects help to control for stable, average differences in household mobility patterns relative to their county, but dynamic, idiosyncratic deviations are not captured. However, the strength of the first stage suggests the instrument effectively captures relevant variation.

While the survey data provides valuable insights for heterogeneity, it may be subject to recall bias, social desirability bias, and the specific timing of the survey relative to the pandemic waves could influence responses.

6 Conclusion

The COVID-19 pandemic, despite its devastating impacts, offered a unique natural experiment to investigate long-standing questions in the economics of addiction. This paper combines large-scale NielsenIQ household purchase data with Advan Patterns mobility data, USAfacts.org COVID-19 and population statistics, and custom survey data to estimate the causal effect of changes in personal cigarette stockpiles on smoking intensity. An Interrupted Time Series analysis documents that the onset of the pandemic in 2020 was associated with statistically significant and economically meaningful average increases in both the quantity of cigarettes purchased per shopping trip and in imputed daily smoking rates by households. Heterogeneity analyses further reveal that these impacts varied across different demographic and psycho-social groups, underscoring the complex ways in which households responded to the pandemic. To isolate the causal effect of interest, an Instrumental Variables strategy was employed, using pandemic-induced changes in county-level shopping mobility as an instrument for household cigarette stock purchases. The analysis, which carefully controls for local COVID-19 case intensity and includes a rich set of household, county, and temporal fixed effects, demonstrates a strong first-stage relationship (F-statistic = 23.6). The resulting IV estimates indicate a significant positive causal effect: an exogenously induced increase in cigarette stock equivalent to one pack (20 cigarettes) is estimated to lead to an increase of approximately one additional cigarette smoked per day by the household. This research provides robust, causally-identified evidence for the "temptation effect" in the context of an addictive good like cigarettes, showing that increased immediate availability from larger personal stockpiles can indeed drive higher consumption. These findings contribute to a deeper understanding of addictive behavior by empirically linking environmental factors (product availability in the home) to consumption intensity, particularly under conditions of significant external shocks. This has important implications for public health policy, cessation program design, and the regulation of addictive goods. Future research could explore the persistence of these pandemic-induced behavioral changes, investigate the role of different types of policy interventions, and examine similar dynamics for other potentially addictive or temptation-inducing

products.

References

Scott R. Baker, R. A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. How does household spending respond to an epidemic? consumption during the 2020 COVID-19 pandemic. *The Review of Asset Pricing Studies*, 10(4):834–862, dec 2020. doi: 10.1093/rapstu/raaa009. URL https://doi.org/10.1093/rapstu/raaa009. 1

Gary S. Becker and Kevin M. Murphy. A theory of rational addiction. 96(4):675–700, aug 1988. ISSN 0022-3808. doi: 10.1086/261558. URL https://www.journals.uchicago.edu/doi/abs/10.1086/261558. Publisher: The University of Chicago Press. 2

Frank J. Chaloupka and Kenneth E. Warner. Chapter 29 the economics of smoking. In *Handbook of Health Economics*, volume 1, pages 1539–1627. Elsevier, 2000. doi: 10.1016/S1574-0064(00)80042-6. URL https://www.sciencedirect.com/science/article/pii/S1574006400800426. 2

Pierre Chandon and Brian Wansink. When are stockpiled products consumed faster? a convenience–salience framework of postpurchase consumption incidence and quantity. 39(3):321–335, aug 2002. ISSN 0022-2437. doi: 10.1509/jmkr.39.3.321.19111. URL https://doi.org/10.1509/jmkr.39.3.321.19111. Publisher: SAGE Publications Inc. 2

Raj Chetty, John N. Friedman, Michael Stepner, and The Opportunity Insights Team. The economic impacts of COVID-19: Evidence from a new public database built using private sector data, jun 2020. URL https://www.nber.org/papers/w27431. 3

Philip DeCicca, Donald Kenkel, and Michael F. Lovenheim. The economics of tobacco regulation: A comprehensive review. 60(3):883–970, sep 2022. ISSN 0022-0515. doi: 10.1257/jel.20201482. URL https://www.aeaweb.org/articles?id=10.1257/jel.20201482. 1, 2

Austan Goolsbee and Chad Syverson. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of Public Economics*, 193:104311, jan 2021. doi: 10.1016/j.jpubeco.2020.104311. URL https://doi.org/10.1016/j.jpubeco.2020.104311. 1

Jonathan Gruber and Botond Köszegi. Is addiction "rational"? theory and evidence. 116(4):1261–1303, 2001. ISSN 0033-5533. URL https://www.jstor.org/stable/2696459. Publisher: Oxford University Press. 2

Faruk Gul and Wolfgang Pesendorfer. Temptation and self-control. 69(6):1403–1435, 2001. ISSN 1468-0262. doi: 10.1111/1468-0262.00252. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00252. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00252. 1, 2

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. 74(6):1637–1673, 2006. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2006.00721.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2006.00721.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2006.00721.x. 2

Hung Ho. Overconsumption and self-control: Evidence from cigarette purchases. SSRN, August 2025. Kilts Center at Chicago Booth Marketing Data Center Paper. Available at SSRN: https://ssrn.com/abstract=5475708. 3

David Laibson. A cue-theory of consumption*. 116(1):81–119, feb 2001. ISSN 0033-5533. doi: 10.1162/003355301556356. URL https://doi.org/10.1162/003355301556356. 1, 2

Asger Lau Andersen, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan. Pandemic, shutdown and consumer spending: Lessons from scandinavian policy responses to COVID-19, may 2020. URL https://ui.adsabs.harvard.edu/abs/2020arXiv200504630L. ADS Bibcode: 2020arXiv200504630L. 3

Ayako Matsuda, Noriko Inakura, and Yuka Sakamoto. Stockpiling behavior during the COVID-19 pandemic: Evidence from a field experiment. Technical report, Consumer Affairs Agency, Government of Japan, 2023. URL https://cir.nii.ac.jp/crid/1010865643882137088. Entry based on user-provided BibTeX. Original source details may vary slightly. CiNii link provided. 3

Brian Wansink. Environmental factors that increase the food intake and consumption volume of unknowing consumers. *Annual Review of Nutrition*, 24:455–479, 2004. doi: 10.1146/annurev.nutr.24.012003.132140. 2

Klaus Wertenbroch. Consumption self-control by rationing purchase quantities of virtue and vice. 17(4): 317–337, nov 1998. ISSN 0732-2399. doi: 10.1287/mksc.17.4.317. URL https://pubsonline.informs.org/doi/10.1287/mksc.17.4.317. Publisher: INFORMS. 2