Hooked on Flavor: Addiction, Present Bias, and the Consequences of E-Cigarette Flavor Policy*

William Brasic†

June 8, 2026

Abstract

Flavored e-cigarettes have captured millions of users who are disproportionately teens and young adults, while bans on them have triggered one of the most highly controversial regulatory debates in recent nicotine public health policy. To evaluate e-cigarette flavor policy, I develop and estimate a dynamic discrete choice model of cigarette and e-cigarette demand with forward-looking, present-biased consumers, featuring nicotine-driven addiction stocks, a flavored habit stock, and heterogeneous flavor preferences that vary with household age composition. Using NielsenIQ scanner data for 2021–2023, I first present reduced-form evidence documenting strong habit persistence, moderate cross-category substitution patterns, and heightened flavored e-cigarette demand among households with younger members. I then estimate the dynamic model via a two-stage pseudo-maximum-likelihood procedure and simulate counterfactual flavor policies under rational and time-inconsistent discounting. Three findings emerge. First, the feared cigarette substitution backfire does not materialize: the dominant response to a ban is market exit rather than switching to cigarettes, driven primarily by strong flavor preferences. Second, enforcement scope is salient: restricting the ban to only FDA-authorized brands captures only 84% of the addiction reduction, with the gap driven by substitution to unauthorized products. Third, a 0.50/mL excise tax achieves 92% of the ban's addiction reduction at less than half the welfare cost. Present bias does not alter this policy ranking, but reduces the true welfare cost of both instruments by converting part of the measured loss into an internality correction.

Keywords: Industrial Organization, Health Economics, Addiction
JEL codes: C61, D12, I18, L66

*I am especially grateful to Ashley Langer and Matthijs Wildenbeest for their feedback, support, and guidance. I also thank Mo Xiao, Christian Cox, Hidehiko Ichimura, Evan Taylor, Dan Herbst, Price Fishback, and participants at The University of Arizona Econometrics Seminar for their insightful comments. Code for the main results of this paper can be found at URL. All errors are my own.

†Department of Economics, The University of Arizona, wbrasic@arizona.edu

I. Introduction

Cigarette smoking kills an estimated seven million people around the world each year, more than any other preventable cause.[1] Economically, cigarette smoking also imposes over $600 billion annually in the United States in healthcare costs and lost productivity.[2] When e-cigarettes emerged commercially in the mid-2000s, they appeared to offer a genuine opening: a nicotine delivery device that bypassed combustion and, if widely adopted, might provide a safer alternative to cigarettes. The potential was real, but so was an unintended consequence. By 2019, more than five million American teenagers were vaping, drawn largely by flavor, and the FDA declared a public health epidemic as nicotine concentration in vapes rose by 106.7%.[3] The policy debate since has centered on two instruments. The first is outright bans on flavored e-cigarettes, enacted by several states beginning in 2019; opponents warn that restricting flavors pushes vapers back to cigarettes, worsening public health rather than improving it.[4] The second is excise taxes, imposed in 33 states and proposed at the federal level through the Tobacco Tax Equity Act of 2023; the open question is whether a price instrument specifically on flavored e-cigarettes can deliver meaningful reductions in nicotine addiction without the welfare cost of an outright ban. Both instruments share a deeper problem. The FDA has authorized only a handful of e-cigarette products, yet by 2023 between 58% and 65% of flavored e-cigarette purchases in the household data used in this paper already flowed through brands with no valid authorization, with the higher share among households with a teen or young adult present, and national retail scanner data suggest that share has since climbed to 86% as of 2024.[5] The fundamental question is not simply which instrument is best in theory; it is which can accomplish anything meaningful given the market it actually has to operate in.

This paper addresses three critical public health policy questions about e-cigarette regulation. First, do flavor bans trigger substitution toward combustible cigarettes, the central empirical objection raised against flavor restrictions? Second, does enforcement scope matter: does a ban that reaches only FDA-authorized products deliver meaningfully different outcomes than a comprehensive ban covering the full market? Third, can e-cigarette excise taxes, already enacted in 33 states and proposed at the federal level through the Tobacco Tax Equity Act of 2023, achieve comparable reductions in nicotine addiction at lower welfare cost? All three questions require accounting for addiction

[1] https://www.who.int/news-room/fact-sheets/detail/tobacco.

[2] https://www.cdc.gov/tobacco/php/data-statistics/economic-trends/index.html

[3] https://truthinitiative.org/research-resources/emerging-tobacco-products/nicotine-content-e-cigarettes-more-doubled-5-years

[4] Cotti et al. (2025); Friedman, Pesko, and Whitacre (2024); Saffer et al. (2025).

[5] https://truthinitiative.org/research-resources/tobacco-industry-marketing/us-retail-sales-data-show-86-e-cigarette-sales-are. The lower shares in the household panel data used here likely reflect both the earlier sample period (through 2023) and the tendency of scanner-based household panels to over-represent mainstream retail channels where FDA-authorized brands have stronger distribution.

dynamics, long-run substitution patterns, and time inconsistency that reduced-form methods, particularly those using quasi-experimental variation in policies implemented during the COVID-19 pandemic, cannot cleanly recover.

Before turning to the structural model, I present reduced-form evidence on nicotine product consumption that motivates both the model's design and the enforcement-gap comparison. I show that households with teens and young adults are significantly more likely to purchase flavored e-cigarettes, that nicotine intake among cigarette users has declined while remaining relatively stable for e-cigarettes, and that category-level transition matrices reveal strong habit persistence alongside moderate cross-category substitution. I also document that the share of flavored e-cigarette purchases involving FDA-authorized brands has declined substantially from roughly 65% in 2021 to 42% in 2023 among households without a teen or young adult, and from 53% to 34% among those with one, suggesting that a flavor ban limited to authorized products would directly affect a shrinking minority of flavored e-cigarette consumption and would disproportionately miss the households policymakers are most concerned about.

To answer the questions central to this paper, I construct and estimate a dynamic structural model grounded in the rational addiction framework of Becker and Murphy (1988). The model features nicotine-based addiction stocks, a flavored habit stock, and heterogeneous flavor preferences that vary with household composition. I then simulate three counterfactual policies: a comprehensive ban on all flavored e-cigarettes, an authorized-only ban targeting only FDA-authorized brands, and per-milliliter excise taxes on flavored e-cigarettes at three rates. The comparison between the first two counterfactuals directly quantifies the enforcement gap: how much of the ban's intended addiction impact is lost because of the majority of e-cigarette sales already flowing through unauthorized channels facing no regulatory exposure. The tax simulations provide an indirect mechanism to reduce flavored e-cigarette consumption. Estimation proceeds via a two-stage pseudo-maximum-likelihood procedure with a full-solution value function routine. The model also allows for quasi-hyperbolic discounting, which I use to identify how the welfare interpretation of the policy changes as consumers become more time-inconsistent.

The main finding from the comprehensive flavor ban is that removing all flavored e-cigarettes from the choice set reduces mean nicotine addiction by 12.2% over the 36-month simulation horizon, with no evidence of a cigarette substitution backfire. The primary behavioral response is exit to the outside option, which rises by 3.9 percentage points over 36 months, while cigarette market share falls by 0.8 percentage points. This result is robust to assumptions about consumer time preferences: across all degrees of present bias considered, the ban delivers addiction reductions within 0.3 percentage points of each other. More present-biased consumers accumulate roughly 130% higher addiction stocks over the simulation horizon, so the absolute addiction decline is larger in levels, but the direction and scale of the effect are stable throughout. The antici-

pated unintended consequence, that flavor bans would push consumers into combustible cigarettes, does not materialize: cigarette market shares fall modestly under the ban, consistent with displaced flavored demand flowing predominantly to market exit rather than to cigarettes.

The response to the ban is heterogeneous across household composition in a pattern favorable to the policy's stated goals. Households with a teen or young adult present exhibit addiction reductions of approximately 15.2% over 36 months, compared to roughly 11.5% among households without youth exposure. This heterogeneity arises because TYA-present households have stronger estimated preferences for flavored products, so the removal of flavors imposes a larger utility reduction on them and pushes a higher fraction toward the outside option. Each percentage point of exit from the nicotine market translates directly into faster addiction stock decay, generating larger long-run addiction reductions for the subgroup targeted by the policy. The comprehensive flavor ban is thus most effective precisely for the demographic motivating it.

To assess the practical reach of a flavor ban under current market conditions, I conduct a second counterfactual that restricts the ban to products sold by brands holding FDA marketing authorization (currently JUUL, Vuse, NJOY, and Logic). These firms face strong compliance incentives: non-compliance with a flavor restriction could jeopardize their existing marketing orders for tobacco and menthol products. By contrast, the unauthorized brands that accounted for between 58% and 65% of flavored e-cigarette purchases in the data already operate outside FDA oversight, and a flavor ban does not meaningfully alter their regulatory exposure. The comparison between the comprehensive and authorized-only bans directly approximates the enforcement gap in U.S. e-cigarette regulation. Among all households, the comprehensive ban reduces mean nicotine addiction by 12.2% over 36 months, while the authorized-only ban reduces it by 10.2%, implying that roughly 16% of the intended addiction impact is lost to the unauthorized market. This attenuation is driven by substitution from authorized flavored products toward unauthorized flavored alternatives that remain available on retail shelves, rather than toward combustible cigarettes. The shortfall is largest for the households the policy is most meant to protect. For TYA households, a comprehensive ban would reduce addiction by 15.2%; a ban limited to authorized brands delivers only 12.4%, a 2.8 percentage point difference. Since teens and young adults already purchase disproportionately from unauthorized brands, an enforcement-limited ban misses its most at-risk demographic the most. The authorized-only ban achieves 84% of the comprehensive ban's addiction reduction while imposing a substantially lower welfare cost: at period 36, the FDA-only ban reduces household welfare by 2.1% relative to the status quo versus 4.9% under the comprehensive ban for TYA households. It also induces consumers to shift to unauthorized firms whose flavored products absorb displaced demand.

Motivated in part by the Tobacco Tax Equity Act of 2023, which proposed a federal

excise tax on all e-cigarettes at a nicotine-equivalent rate of cigarettes, I also simulate per-milliliter excise taxes on flavored e-cigarette liquid of \$0.10, \$0.25, and \$0.50 per milliliter. The \$0.50/mL rate reduces mean addiction by 11.2%, approaching the comprehensive ban's impact, at less than half the welfare cost of an outright ban. Taxes thus dominate outright bans on welfare-efficiency grounds: the \$0.50/mL tax achieves 92% of the comprehensive ban's addiction reduction at less than half the welfare cost while also generating government revenue. Notably, \$0.50/mL is far below the implied per-milliliter burden of the proposed Tobacco Tax Equity Act of 2023, which would translate to approximately \$2.78/mL for e-cigarettes at the conventional rate of 50 mg of nicotine per milliliter of e-liquid, suggesting that even modest volume-based taxes deliver substantial public health gains well short of the federal proposal.[6] The one exception is among the most present-biased households with teens or young adults: there the \$0.50/mL tax falls just short of the ban (14.9% vs. 15.4% addiction reduction), suggesting that when time inconsistency is severe and youth exposure is present, a product set restriction may be more effective over the long-run.

This paper builds on two strands of literature: dynamic models of rational addiction and empirical evaluation of e-cigarette flavor policy. On the addiction side, Becker and Murphy (1988) first proposed that even for addictive goods, consumers behave as forward-looking utility maximizers. Subsequent work has found evidence consistent with this framework (Becker, Grossman, and Murphy (1994); Baltagi and Griffin (2002); Bask and Melkersson (2004); Piccoli and Tiezzi (2021); Deng, Zheng, and Butler (2023)), while Chaloupka (1991) documents elasticity heterogeneity by age, a critical demographic in my paper. Dynamic structural applications of this model include Gordon and Sun (2015), who model addiction under cigarette tax policies and show that failing to account for addiction leads to underestimated price elasticities, and Chen and Rao (2020), who develop a dynamic model of joint cigarette and e-cigarette demand. On the flavor ban side, Tsai et al. (2016) find that availability of flavors is the second most common reason for e-cigarette use among middle and high school students. Chen and Reinhardt (2025) construct a static demand model finding no evidence of spillovers into cigarettes following a simulated flavor ban. By contrast, reduced-form evidence suggests that such bans may increase cigarette consumption as a substitution effect (Cotti et al. (2025); Friedman, Pesko, and Whitacre (2024); Saffer et al. (2025)), though most of this evidence comes from policies implemented in early 2020 contemporaneously with the COVID-19 epidemic. My approach addresses this by adopting a strategy analogous to Petrin (2002), who used a structural model to evaluate the welfare effects of the minivan by simulating a counterfactual world without it based on post-entry data. In a similar spirit, I use my structural model to simulate a counterfactual in which flavored e-cigarettes are removed from the choice set, allowing for a cleaner assessment of the

[6] The bill sets a rate of \$100.66 per 1,810 mg of nicotine, implying $\$100.66/1,810 \approx \$0.0556$ per mg. Multiplying by a nicotine concentration of 50 mg/mL gives $\$0.0556 \times 50 \approx \$2.78$ per mL for a typical high-nicotine disposable.

policy's long-run impact on addiction and substitution behavior.

My paper extends this literature in multiple ways. First, it resolves three open empirical and policy questions that the existing literature either has limited evidence on or has left entirely unaddressed: whether flavor bans trigger cigarette substitution, how much of a ban's intended impact survives incomplete enforcement, and whether a per-milliliter excise tax is a viable alternative. Prior work on substitution relies on reduced-form variation from policies enacted contemporaneously with COVID-19; prior work on bans evaluates them as if they bind on the full market rather than only on the authorized segment; and no study has structurally compared a flavored e-cigarette excise tax to an outright ban. By modeling addiction dynamics, heterogeneous flavor preferences, and a consumer proclivity towards products from compliant and non-compliant firms, this paper provides a unified framework for answering all three questions and directly informs the active federal debate over the FDA authorization process and the Tobacco Tax Equity Act of 2023. On the methodological side, it advances the dynamic structural literature on addiction in three respects. A flavored habit stock is added alongside the standard nicotine addiction stock, which is necessary for capturing dynamic substitution patterns following a flavor ban or tax. Addiction is driven by nicotine intake rather than consumption quantities, following Adda and Cornaglia (2006), which allows cigarettes and e-cigarettes to differ in their addiction consequences per unit consumed rather than treating all purchases symmetrically.

Moreover, I incorporate quasi-hyperbolic discounting (Laibson (1997)) to assess how present-biased preferences reshape the welfare interpretation of flavor policy. Under rational discounting, the ban's measured welfare cost is straightforward: consumers are losing access to products they chose freely and valued. However, under time-inconsistent preferences part of that cost reflects an internality correction, since present-biased consumers over-consume flavored e-cigarettes relative to their own long-run preferences, partially offsetting the social loss, consistent with the framework of Gruber and Köszegi (2001). This present-bias analysis serves as a welfare sensitivity check: policy conclusions are qualitatively robust across all discounting assumptions, but how much of the welfare cost counts as genuine harm depends on whether purchase behavior reflects true preferences or a self-control failure. The data are suggestive on this point: 68% of adult smokers report wanting to quit and over half attempt to in any given year, yet fewer than 8% succeed (Babb et al. (2017)); among daily youth JUUL users, over 60% report having tried to quit and failed (Kechter et al. (2021)). A fully rational consumer who wanted to quit would quit. The gap between stated intentions and actual behavior is what present bias predicts. I model consumers as sophisticated so that they know they are present-biased and plan accordingly, following O'Donoghue and Rabin (1999).

The remainder of this paper is organized as follows. Section II. provides an overview of the history and current landscape of the nicotine product market. Section III. describes the data sources used for the analyses. Section IV. reports descriptive and

reduced-form results that motivate the research questions and dynamic addiction model. Section V. details the structural model, while Section VI. discusses identification and estimation. Section VII. examines the estimated parameters, their significance, and model fit. Section VIII. presents the counterfactual results across all policies considered along with comparison of their relative efficacy. Finally, Section IX. concludes with final remarks on the findings, limitations, and areas for further research. Section X. provides supplementary results, additional theoretical details, and pseudocode for the estimation procedure.

II. Industry Background

This section describes the U.S. cigarette and e-cigarette markets, the FDA's authorization framework, and the state-level flavor ban landscape that contextualizes the counterfactual analysis. Two features of this market are central to interpreting the structural results. First, cigarettes and e-cigarettes are linked by nicotine demand: consumers deterred from e-cigarettes by a flavor ban may substitute toward cigarettes rather than exit the market, with important implications for public health. Second, the majority of the e-cigarette market currently operates outside the FDA's authorization framework, so a ban covering only authorized products leaves most of the market untouched. Both features shape the paper's key counterfactual finding that a comprehensive flavor ban produces dramatically larger reductions in addiction than an FDA-only ban. Figure 1 traces the aggregate trend over the sample period: cigarette volumes declining while e-cigarette purchases rise, a pattern that flavor bans could reinforce or, the substitution hypothesis warns, mitigate.

Figure 1. Median Monthly Consumption of Packs of Cigarettes (Panel (A)) and Milliliters of E-Liquid (Panel (B)).

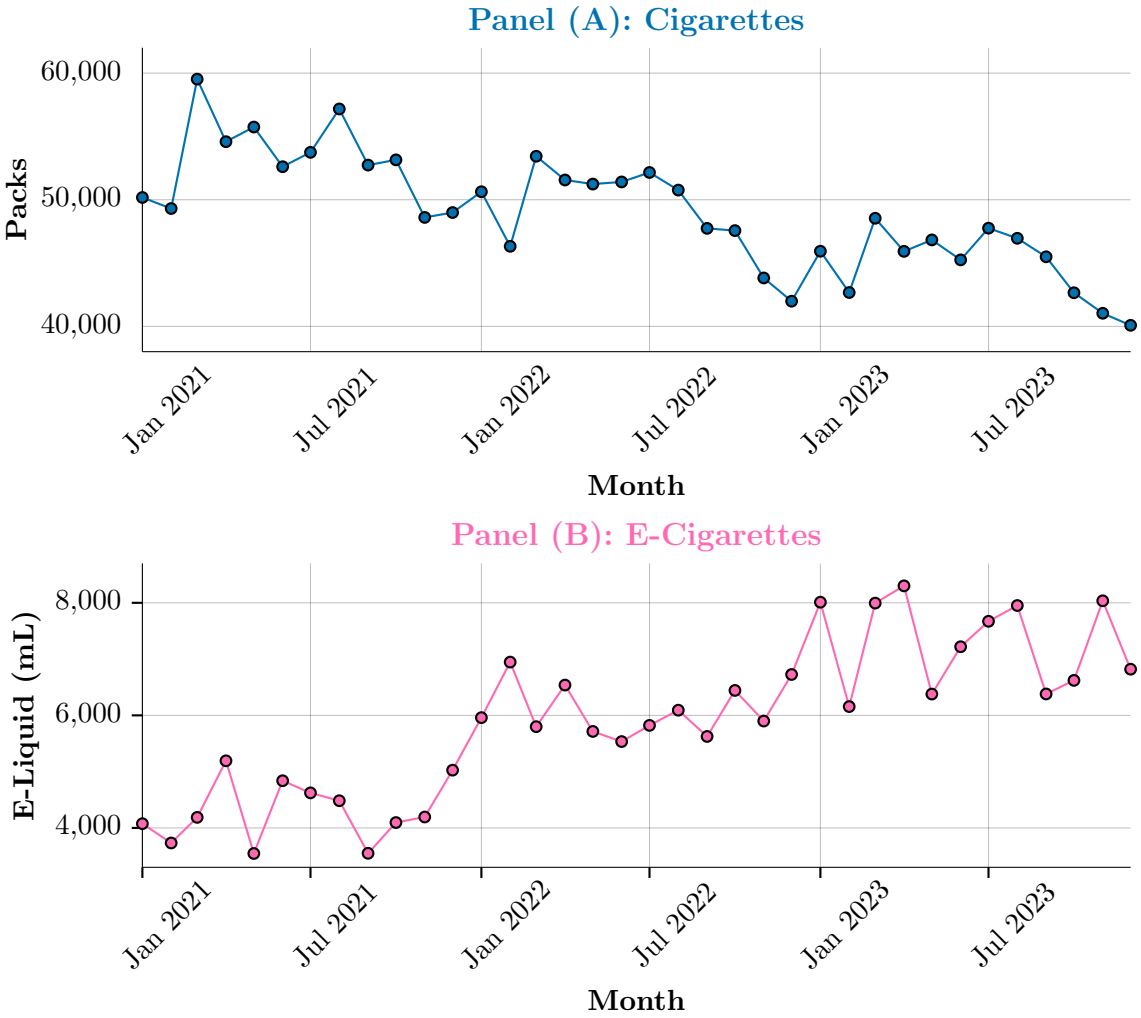

This figure is based on data of those who purchased cigarettes (Panel (A)) and e-cigarettes (Panel (B)), implying this is analysis of the intensive margin. I filter out households not appearing in all years to form a balanced panel to avoid confounding the effect of an increase (decrease) in household purchase frequency with the entry/exit of heavy (light) cigarette/e-cigarette users.

II.1. *Cigarettes: Decline, Market Structure, and Addiction-Driven Persistence*

Combustible cigarettes have been the dominant nicotine delivery product in the United States for over a century. Adult smoking prevalence peaked above 42% in the mid-1960s and has since declined following the 1964 Surgeon General's report linking smoking to lung cancer, successive federal and state excise tax increases, advertising restrictions, and growing public health awareness.[7] Despite this decline, an estimated seven million people die of smoking each year,[8] making cigarettes the leading cause

[7] Adult cigarette smoking has declined from 42.6% in 1965 to 11.6% in 2022 (https://www.lung.org/research/trends-in-lung-disease/tobacco-trends-brief/overall-smoking-trends).

[8] https://www.who.int/news-room/fact-sheets/detail/tobacco

of preventable death in the world. Moreover, cigarette smoking also imposes over $600 billion annually in the United States in healthcare costs and lost productivity.[9] The market is highly concentrated: three manufacturers, Altria (Philip Morris USA), Reynolds American (R.J. Reynolds), and ITG Brands, account for roughly 95% of domestic retail sales. Retail prices vary substantially across states, primarily driven by state excise taxes that currently range from under $0.20 per pack to over $4.00 per pack.

The persistence of cigarette demand despite decades of well-publicized health risks reflects the addictive nature of nicotine. Price elasticities of cigarette demand are consistently estimated to be modest in magnitude and substantially smaller than most consumer goods, reflecting addiction-driven inertia. This inertia makes cigarettes the natural substitution target when e-cigarette alternatives are restricted: consumers whose nicotine dependence is high enough to deter market exit, but who are deterred from switching to unflavored products by flavor-specific habit formation, may turn to cigarettes. The public health concern is that a flavor ban could reduce e-cigarette use while simultaneously increasing cigarette smoking, resulting in net harm. Resolving this empirically requires a structural model that separately identifies addiction dynamics, flavor-specific inertia, and the relative utility of each product category, the core contribution of Section V..

II.2. *E-Cigarettes: Devices, Market Structure, Flavors, and Taxes*

E-cigarettes (also called vapes or electronic nicotine delivery systems) heat a nicotine-containing liquid into an aerosol for inhalation, replacing combustion with vaporization (Figure 2, Panel (A)). The market spans three main device formats: pod-based closed systems, open refillable tanks, and single-use disposables (Figure 2, Panel (B)). Modern closed systems and disposables typically use nicotine salts, a chemically modified form of nicotine that tolerates high concentrations without throat irritation. Popular disposables advertise nicotine concentrations of 50 mg/mL, so a standard 2 mL device contains approximately 100 mg of nicotine, roughly half of which is absorbed during use, delivering more absorbed nicotine than several packs of cigarettes in a single device. Between 2013 and 2018, the total amount of nicotine delivered via e-cigarettes in the U.S. increased by approximately 106.7%, driven by a shift toward larger, higher-concentration disposables.[10] The global e-cigarette market size reached roughly $45 billion in 2025.[11] Three FDA-authorized brands dominate the legal market: Vuse (Reynolds American), JUUL (JUUL Labs), and NJOY (acquired by Altria in 2023) together account for the vast majority of authorized market share as of early 2024, with Vuse alone holding

[9] https://www.cdc.gov/tobacco/php/data-statistics/economic-trends/index.html

[10] https://truthinitiative.org/research-resources/emerging-tobacco-products/nicotine-content-e-cigarettes-more-doubled-5-years

[11] https://www.grandviewresearch.com/industry-analysis/e-cigarette-vaping-market

roughly 42%.[12] To legally market an e-cigarette in the U.S., a manufacturer must obtain a Premarket Tobacco Application (PMTA) marketing order from the FDA by demonstrating the product is "appropriate for the protection of public health," a bar that most products currently on the market have not cleared.

Figure 2. E-Cigarettes: How They Work and Their Evolution Over Time.

Panel (A): Anatomy of an E-Cigarette

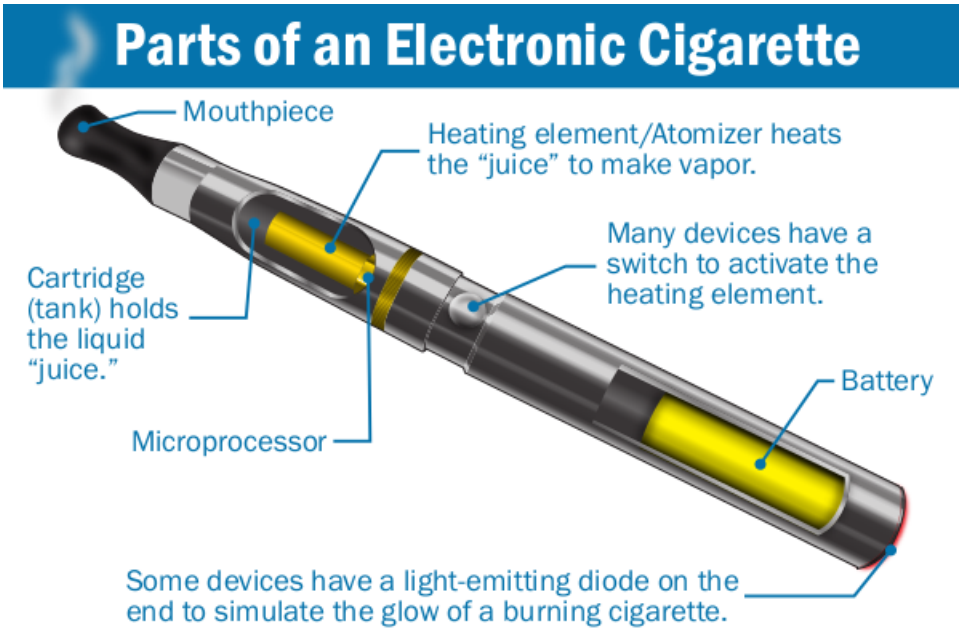

Panel (B): E-Cigarette Variations

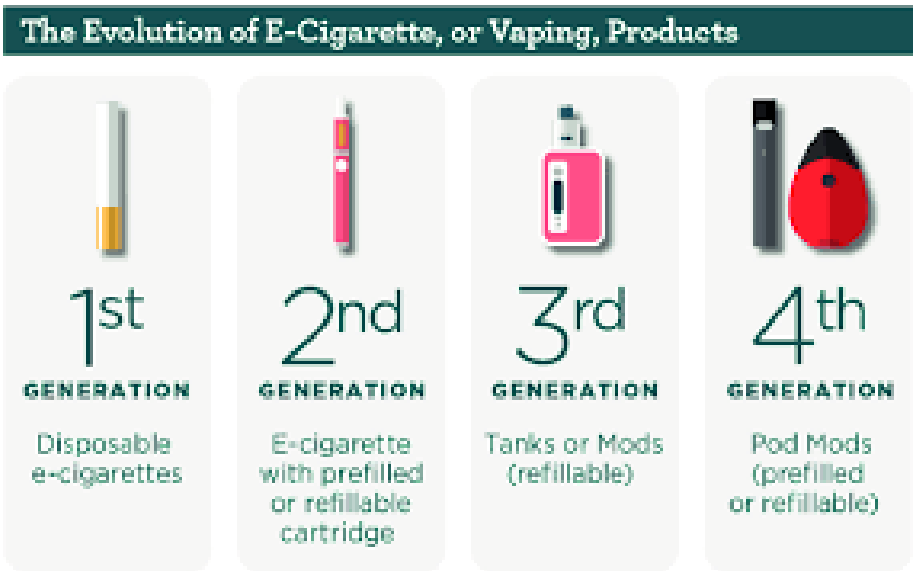

E-cigarettes are sold in a wide range of non-tobacco flavors, including fruit, candy,

[12] JUUL held 75% market share at its peak in May 2019 before FDA enforcement actions and the rise of unauthorized disposable brands sharply eroded its position.

dessert, mint, and menthol varieties. Flavored products are disproportionately popular among younger users: non-tobacco flavors mask the harshness of nicotine, lower the perceived barrier to initiation, and, in combination with high nicotine salt concentrations, accelerate dependence among adolescents. According to the 2023 National Youth Tobacco Survey (NYTS), 10% of U.S. middle and high school students (approximately 2.8 million youth) reported current e-cigarette use, and among those users, 89.4% reported using a flavored product.[13] This flavor appeal is a primary motivation for the demographic heterogeneity in this paper: households with a teen or young adult (TYA) member exhibit markedly different e-cigarette purchasing patterns than non-TYA households, and the structural model captures this through TYA-specific flavor preference parameters.

Beginning around 2019–2020, several states enacted statewide bans on the retail sale of most flavored e-cigarettes, largely filling a regulatory void left by federal inaction. Massachusetts enacted the first comprehensive state ban in November 2019, notable for its breadth: it covered not only flavored e-cigarettes but also menthol cigarettes and flavored cigars, making it the most far-reaching state tobacco restriction in the country at the time. New York, New Jersey, Rhode Island, and Utah followed in 2020, with California following suit in November 2022 after voters upheld the state's flavor ban via Proposition 31. By 2023, over 375 U.S. localities had adopted permanent flavor restrictions of some kind. The laws vary in scope: some cover only non-tobacco-flavored e-cigarettes, while others extend to menthol cigarettes and flavored cigars. All operate at the retail level, restricting licensed point-of-sale transactions but leaving significant gaps in enforcement. Online sales, cross-border purchases from neighboring states, and direct-to-consumer shipments from overseas manufacturers (particularly Chinese producers supplying the large unauthorized disposable vape market) fall largely outside these restrictions. The FDA has repeatedly flagged the importation of unauthorized flavored e-cigarettes from China as a primary enforcement challenge, as these products enter the U.S. market without premarket review and account for a substantial share of flavored e-cigarette availability despite being illegal to sell domestically. These restrictions are uniformly framed as youth-protection measures, reflecting evidence that flavored products are disproportionately used by adolescents. The tobacco and vaping industries have challenged several of these laws in court, with mixed results, and the patchwork of state and local rules has created significant geographic variation in product availability.

A large share of the U.S. e-cigarette market operates outside the FDA authorization framework. As of January 2026, the FDA has authorized only 39 e-cigarette products across all flavor categories, nearly all tobacco-flavored, with the first non-tobacco flavor authorization (four NJOY menthol products) granted in June 2024.[14] Despite this

[13] https://www.fda.gov/news-events/press-announcements/national-survey-shows-drop-e-cigarette-use-among-high-school-students

[14] https://www.fda.gov/tobacco-products/market-and-distribute-tobacco-product/tobacco

narrow authorized set, unauthorized products dominate the retail market. Truth Initiative's 2024 analysis of retail scanner data finds that products lacking a valid PMTA marketing order account for approximately 86% of U.S. e-cigarette dollar sales.[15] This segment is dominated by heavily flavored, high-nicotine disposables manufactured primarily in China, including brands such as Elf Bar, Breeze Smoke, Flum, and Geek Bar, sold through the same retail channels as authorized products. Among youth, 56% of e-cigarette users reported using Elf Bar in 2023, making it the most prevalent brand among adolescents despite having no FDA authorization.[16] This market structure creates what I call the enforcement gap: a flavor ban covering only FDA-authorized products leaves unauthorized flavored substitutes freely available, so consumers need not change their behavior. Only a ban that extends to all flavored products, authorized or not, can substantially reduce flavored e-cigarette consumption and its downstream effects on nicotine addiction. I quantify this gap in Section VIII.1. by comparing simulated outcomes under a comprehensive ban against an FDA-only ban, illustrating how the health impact of flavor regulation depends critically on its scope.

As an alternative policy instrument, excise taxes on e-cigarette liquid have gained traction at both the state and federal levels. As of September 2024, 33 states plus the District of Columbia tax e-cigarette products in some form, either as a percentage of wholesale price or as a rate per milliliter of e-liquid.[17] At the federal level, no e-cigarette excise tax currently exists. Unlike cigarettes, which face a federal rate of $1.01 per pack in addition to state taxes, e-cigarettes have until now been taxed only at the state level, leaving a large portion of the market subject to no excise tax at all in states without their own e-cigarette levy. The most recent federal proposal to close this gap is the Tobacco Tax Equity Act of 2023, introduced by Senator Dick Durbin and eight cosponsors. The bill would impose a federal excise tax on all nicotine-containing products at a rate equivalent to the cigarette tax: $100.66 per 1,810 milligrams of nicotine, or approximately $0.056 per milligram of nicotine.[18]

-products-marketing-orders

[15] https://truthinitiative.org/research-resources/tobacco-industry-marketing/us-retail-sales-data-show-86-e-cigarette-sales-are

[16] https://www.fda.gov/news-events/press-announcements/national-survey-shows-drop-e-cigarette-use-among-high-school-students

[17] Tax structures vary considerably: Minnesota imposes 95% of wholesale price (the highest ad valorem rate in the country), Vermont 92%, Massachusetts 75%; Connecticut imposes $0.40 per milliliter (the highest per-volume rate) on closed systems (pod-based or disposable), New Hampshire $0.30 per milliliter, and several states as little as $0.05 per milliliter. (https://taxfoundation.org/data/all/state/vaping-taxes/)

[18] The per-milligram rate is $\$100.66/1,810 \approx \$0.0556$ per mg of nicotine. To convert to a per-mL rate, multiply by the product's nicotine concentration: at 50 mg/mL, typical of popular disposable brands such as Elf Bar, this gives $\$0.0556 \times 50 \approx \$2.78$ per milliliter, more than six times Connecticut's $0.40/mL rate and far above the rates I simulate. A lower-concentration open-system liquid at 25 mg/mL would face only about $1.39/mL under the same bill, illustrating how the nicotine-content-based structure hits high-concentration disposables far harder than lower-concentration products. The bill was referred to the Senate Finance Committee and did not pass.

III. Data

III.1. *Household Scanner Data*

The core results of this paper use the NielsenIQ Consumer Panel.[19] NielsenIQ is a global data analytics company that provides detailed insights into consumer purchasing behavior by collecting and analyzing retail scanner and household panel data across a wide range of product categories and geographic markets.

The Consumer Panel Data represents a longitudinal panel of approximately 60,000 U.S. households from 2004 onward who continually provide information to NielsenIQ about their households, products they buy, as well as when and where they make purchases. Households in the NielsenIQ Consumer Panel are selected using a stratified and proportionate (proportional to their share of the total population within each strata) sampling design intended to be nationally representative of U.S. households based on demographics such as age, income, household size, and geographic region. A variety of online vendors are used to provide household email information and random site invitation recruitment to recruit new panelists online.[20] The panelists are demographically balanced and geographically dispersed with information on each panelist's state, county, and zip code. NielsenIQ samples all states and major markets (except Alaska and Hawaii). Some panelists stay on the panel for several years, while others may join or drop off each year. Demographic information is recorded for the entire household and the head of household. Demographic variables include, but are not limited to, household size, income, age, presence and age of children, employment, education, marital status, occupation, type of residence, and race. NielsenIQ has a comprehensive program of dropping and replacing panelists that do not perform to minimum reporting standards. Currently, NielsenIQ retains about 80% of its active panel each year.

NielsenIQ Homescan panelists use in-home scanners or a mobile app to record all purchases intended for personal, in-home consumption. The panel captures products across all NielsenIQ-tracked food and non-food categories from all retail outlets in U.S. markets, including online purchases from retailers such as Instacart. NielsenIQ estimates that approximately 30% of total household consumption is represented by the consumer panel categories. For each shopping trip, the data contain summary information such as the household identifier, trip date, retailer and store codes, store ZIP code, and total expenditure. Within each trip, detailed transaction-level information is provided for every product purchased, including the universal product code (UPC), quantity, price,

[19] Researcher(s) own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. See https://www.chicagobooth.edu/research/kilts/research-data/nielseniq, https://nielseniq.com/global/en/info/consumer-panel-data/, and https://nielseniq.com/global/en/products/retail-measurement-services-rms/ for further details.

[20] Panelist recruitment is only done online.

and any associated deals or coupons.

I use NielsenIQ Consumer Panel data covering the years 2021 – 2023, excluding 2020 due to distortions arising from the COVID-19 pandemic, with 2023 being the most recent available year. Approximately 25% of households are observed in all three years. The analysis focuses on households that made at least one tobacco-related purchase during this period, thereby capturing behavior along the intensive margin of tobacco consumption. Tobacco products are defined as belonging to one of three categories: combustible cigarettes and e-cigarettes in either original or flavored form. I exclude e-cigarette purchases that are not cartridges, e-liquid refills, or disposables, as these three categories account for roughly 95% of all e-cigarette transactions in the panel between 2021 and 2023, inclusive. Chargers, batteries, and starter kits comprise the remaining share; starter kits are typically one-time purchases and often do not include e-liquid, making flavor data unreliable.

The final sample consists of 3,094 unique UPCs: 1,662 for cigarettes and 1,432 for e-cigarettes. While the NielsenIQ Consumer Panel data include product descriptions such as brand information, they typically lack details on whether an e-cigarette product is flavored, or the amount of e-liquid contained in milliliters and its nicotine concentration. These characteristics are central both to the nature of the products and to the policy analyses undertaken in this paper. To address this limitation, I used the UPCs for e-cigarettes to obtain this missing information.[21] Through this process, I obtain information on whether each product is flavored, its nicotine concentration, and the corresponding quantity of e-liquid for the given purchase.[22] Nicotine strength was not given for any product in the NielsenIQ data, so I use UPC lookups to retrieve nicotine concentrations in mg/mL for all 1,432 e-cigarette products. Finally, for any e-cigarette product with an invalid UPC in online databases, I use the product description provided by NielsenIQ to manually search for and verify the missing information using Google.[23]

I also classify each e-cigarette brand as FDA-authorized (JUUL, Vuse, NJOY, Logic) or unauthorized based on whether the brand held a valid PMTA marketing order as of 2026. The authorized-versus-unauthorized distinction reflects brand-level regulation rather than product-level approval.[24] Within the data, unauthorized brands account for

[21] I used https://go-upc.com/, which markets itself as "The World's Largest Barcode Database." Because I assume a constant nicotine content per cigarette and do not distinguish between menthol vs. non-menthol in the data, I did not need to look up cigarette UPCs.

[22] For the few e-cigarette products for which the e-liquid quantity (in mL) was unavailable, I use the reported maximum number of puffs to infer the amount of e-liquid, assuming approximately 300 puffs per mL.

[23] Certain purchases display implausible prices, such as instances where a household pays nearly $1,000 for a single pack of cigarettes or reports paying nothing for a carton (10 packs). To address these outliers, I adjust per-unit cigarette prices by replacing values in the bottom and top 0.5% of the per-pack price distribution with the respective median per-unit prices within those tails. If the median value remains $0.00 within the lower tail, I substitute it with the maximum per-pack price observed in the lower 0.05% of the distribution to ensure nonzero prices. I apply a similar procedure to e-cigarettes, using the bottom and top 1% of the per-mL e-liquid price distribution to replace extreme or erroneous values. The wider trimming range reflects the smaller number of e-cigarette purchases relative to cigarette purchases.

[24] A finer product-level classification would require matching each UPC to its specific PMTA submis-

between 58% and 65% of flavored e-cigarette purchases, a lower share than national retail scanner estimates as of 2024 (approximately 86%), likely because the HMS panel skews toward older, established purchasers with stronger loyalty to legacy authorized brands, and because NielsenIQ's UPC tracking infrastructure more reliably captures products sold through established retail channels, whereas many unauthorized disposables enter the market through informal distributors or online. As a result, counterfactual estimates of the enforcement gap in this paper are best interpreted as a lower bound on the addiction reduction lost to incomplete enforcement.

Since the NielsenIQ Consumer Panel tracks household-level purchases, the estimates in this paper primarily reflect the behavior of adult household purchasers. Purchases made independently by youth using their own money, particularly through informal channels, are not fully captured in this dataset. This limitation is partially offset by the panel's coverage of household retail activity: a substantial share of youth e-cigarette consumption passes through retail channels the panel is likely to record. A 2023 study finds that 43.1% of 15–20-year-old e-cigarette users obtained their devices from retail sources, rising to 51.4% among 18–20-year-olds and to 50.0% for disposable devices specifically, with approximately 60% of underage retail purchases involving fruit, candy, dessert, or menthol/mint/ice flavors.[25] The remaining 56.9% obtained e-cigarettes through social sources, including friends and family members who may themselves be panel members making household purchases. Rather than tracking youth consumption directly, I study household-level heterogeneity by the presence of teens or young adults (TYA), which captures both household purchases involving youth and the influence of youth preferences on household purchasing decisions. Households with TYA present show systematically stronger preferences for flavored and unauthorized products in the data, consistent with youth driving demand toward the segment most relevant to the enforcement gap counterfactual. My estimates of the TYA-specific enforcement gap are therefore best interpreted as a lower bound: by missing purchases youth make independently outside the household, the panel understates TYA households' true exposure to unauthorized flavored products, which in turn understates the policy responsiveness of this subgroup relrative to what survey-based sources such as the NYTS or the Population Assessment of Tobacco and Health (PATH) study would suggest.[26]

sion and authorization status, which is not systematically available in any public database. Brand-level classification is a reasonable proxy because compliance incentives operate primarily at the brand level: a firm that holds marketing orders for any of its products faces reputational and regulatory exposure that shapes its behavior across its entire product line, while brands with no authorized products face no such incentives regardless of which specific variants they sell. The counterfactual ban on FDA-authorized brands is therefore best understood as a ban on the product lines of compliant firms rather than a product-by-product restriction.

[25] https://truthinitiative.org/research-resources/tobacco-industry-marketing/new-study-43-underage-e-cigarette-users-report.

[26] I choose not to use these datasets for several reasons: (1) they have already been employed to evaluate similar policies (Cotti et al. (2025), Saffer et al. (2025)); (2) the NYTS does not follow the same respondents over time; and (3) while PATH is longitudinal, it only includes behavioral questions such as "Have you smoked in the last month?" and lacks information on prices paid or specific product

IV. Empirical Motivation

This section presents a series of descriptive and reduced-form exercises that motivate the structural model and establish the key empirical patterns the model is designed to explain. I begin with summary statistics comparing demographic characteristics across household types to show that tobacco-purchasing and non-purchasing households are broadly balanced on observables. I then examine trends in nicotine intake using month fixed effects regressions to distinguish intensive from extensive margin dynamics. Next, I document that households with a teen or young adult present purchase flavored e-cigarettes at substantially higher rates, both unconditionally and in OLS regressions controlling for demographics, state, and month fixed effects. I then present month-to-month product transition matrices and an abstention spell analysis to characterize habit persistence and show time inconsistency is present in the data, respectively. I close with evidence on the declining share of FDA-authorized brands among flavored e-cigarette purchases, which motivates the enforcement gap counterfactual.

IV.1. *Summary Statistics and Stylized Facts*

Table A.1 provides definitions of the key variables used in this section, while Table 1 presents summary statistics of household demographics across selected subgroups. Panel (a) compares households that purchase at least one inside good (Inside Good Households) with those that never purchase an inside good over the panel period (Outside Good Households). The results indicate that households purchasing tobacco-related products and those abstaining from such products are broadly similar across observable characteristics, aside from non-tobacco households tending to have higher incomes and educational attainment. These patterns are also consistent with the averages reported in the final column using 2021 – 2023 U.S. Census data. Given the close demographic alignment between tobacco-using and non-tobacco-using households, as well as their comparability to the broader population, the intensive margin results discussed in this paper are likely to generalize to the extensive margin and to U.S. households more broadly.

characteristics such as nicotine concentration, which are central to the analyses in this paper.

Table 1. Household Descriptive Statistics.

Panel (A)	Inside Good Households		Outside Good Households		2021 – 2023 Census	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Household Income	$51,185	$29,929	$64,241	$29,423	$66,655	$33,384
HoH Age	56	12	58	14	55	17
HoH w/ College Degree	0.35	0.47	0.58	0.48	0.36	0.48
White Households	0.80	0.40	0.77	0.42	0.64	0.48
Married Households	0.53	0.49	0.62	0.48	0.45	0.50
Child Present	0.14	0.33	0.14	0.34	0.17	0.38
Teen Present	0.13	0.32	0.13	0.31	0.13	0.33
Young Adult Present	0.12	0.31	0.11	0.29	0.15	0.35
Number of Households	9,675		62,208		4,464,983	

Panel (B)	Cig. Households		E-Cig. Households		Cig. & E-Cig. Households	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Household Income	$51,212	$29,868	$55,173	$30,218	$48,291	$30,112
HoH Age	56.14	12.18	53.04	12.07	50.84	11.94
HoH w/ College Degree	0.35	0.47	0.42	0.49	0.33	0.45
White Households	0.79	0.40	0.80	0.40	0.84	0.36
Married Households	0.54	0.49	0.49	0.50	0.49	0.49
Child Present	0.13	0.33	0.14	0.34	0.17	0.36
Teen Present	0.12	0.32	0.15	0.35	0.17	0.36
Young Adult Present	0.11	0.30	0.16	0.36	0.16	0.34
Number of Households	8,382		511		782	

Census data for 2021 – 2023 is the American Community Survey (ACS) Public Use Micro Data (PUMS) gathered from https://www.census.gov/programs-surveys/acs/microdata/access.html, which contains information on individual people and housing units. "Inside Good Households" are defined as those households who made at least one cigarette or e-cigarette purchase over the 2021 – 2023 time period. "Cig. Households" are those households who made at least one cigarette purchase during 2021 – 2023 and never purchased e-cigarettes during that entire period. "E-Cig. Households" are those households who made at least one e-cigarette purchase during 2021 – 2023 and never purchased cigarettes during that entire period. "Cig. & E-Cig. Households" are those households who made at least one cigarette purchase and at least one e-cigarette purchase at any point during 2021 – 2023.

Panel (b) of Table 1 provides additional breakdowns by product category, distinguishing between households that purchase only cigarettes (Cig. Households), only e-cigarettes (E-Cig. Households), and both (Cig. & E-Cig. Households). Again, these groups appear balanced across observable characteristics, with the majority of households being cigarette-only consumers. Notably, among e-cigarette–purchasing households, roughly 44% include at least one teen or young adult, a substantially higher share compared to households that purchase only cigarettes.

Figure 3. Month Fixed Effects Relative to January 2021 for Regression of the Natural Log of Nicotine Purchased for Cigarettes (Panel (A)) and E-Cigarettes (Panel (B)).

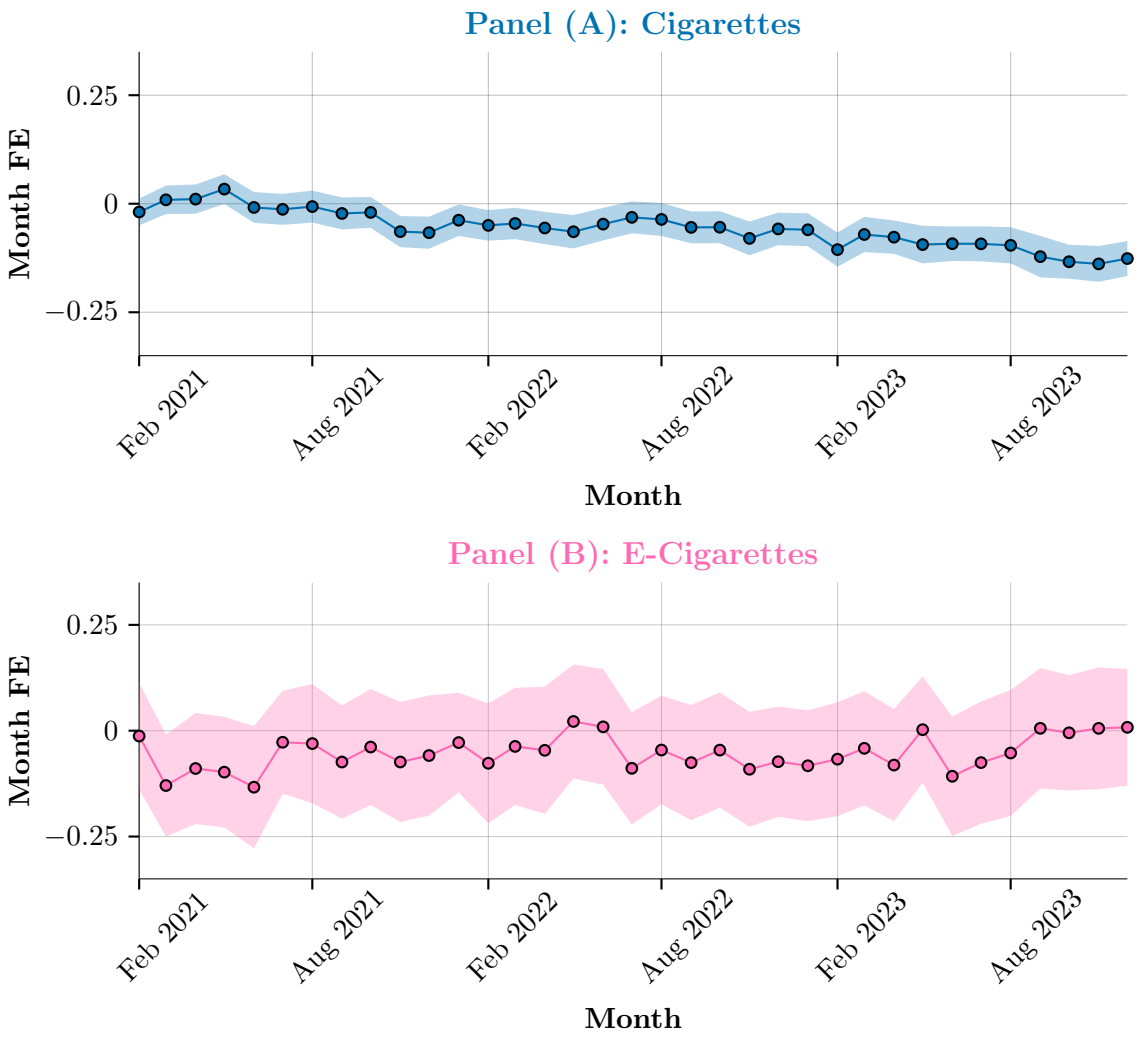

This figure is based on data of those who purchased cigarettes (Panel (A)) and e-cigarettes (Panel (B)), implying this is analysis of the intensive margin. I filter out households not appearing in all years to form a balanced panel to avoid confounding the effect of an increase (decrease) in household purchase frequency with the entry/exit of heavy (light) cigarette/e-cigarette users. Other covariates include household fixed effects, total expenditures, household income, and the number of individuals in the household. Shaded areas correspond to 95% confidence bands. Cluster-robust standard errors are at the household level.

While Figure 1 shows a decline in nicotine intake from cigarettes accompanied by a rise in nicotine intake from e-cigarettes, it remains unclear whose behavior underlies these trends. Are new consumers entering the e-cigarette market, or are existing users increasing their consumption intensity? Furthermore, if new users are entering, could this shift be driven by the decline in inflation-adjusted prices observed in Figure A3? To distinguish between movements along the intensive and extensive margins, while accounting for changes in inflation-adjusted expenditures, I estimate regressions of the natural log of nicotine intake for both cigarette and e-cigarette purchases. These models

control for real prices, a rich set of household characteristics, and include both household and month fixed effects. The estimated month fixed effects are presented in Figure 3.

Two key insights emerge. First, conditional on purchasing, Panel (a) shows cigarette users within households exhibit declining nicotine intake over time by roughly 13% relative to January 2021, consistent with reductions along the intensive margin.[27] Given the median number cigarettes purchased in a month is 200 (10 packs), this implies a within-household reduction of roughly 26 cigarettes (about 1.3 packs) monthly over the 2021 – 2023 time horizon. Second, Panel (b) illustrates that e-cigarette users show little change in conditional nicotine intake, implying that the overall rise in e-cigarette consumption is driven primarily by the extensive margin: new consumers entering the market rather than existing users increasing their consumption intensity.

Figure 4 presents purchase frequencies of e-cigarette purchases categorized as original (non-flavored) or flavored. This figure reveals an important pattern: households with at least one teen (aged 13 – 19, inclusive) or young adult (aged 20 – 25, inclusive) purchase flavored e-cigarettes at roughly four times the rate they purchase non-flavored ones, compared to a two-to-one ratio among households without a teen or young adult, where the former category is subject to flavor bans in several U.S. states.[28] Given that most states (forty-one since 2023) adhere to the 2019 federal restriction requiring individuals to be at least 21 years old to purchase e-cigarettes, this finding suggests that teens and young adults may either be obtaining products through informal or illegal channels. A 2023 study finds that 43.1% of 15–20-year-old e-cigarette users obtained their devices from retail sources, rising to 51.4% among 18–20-year-olds and to 50.0% for disposable devices specifically, with approximately 60% of underage retail purchases involving fruit, candy, dessert, or menthol/mint/ice flavors.[29]

[27] I assume a constant rate of nicotine per cigarette of 12mg implying Panel (a) is the same as a regression on the log of cigarette consumption.

[28] It is important to note that these estimates likely understate total teen and young adult e-cigarette consumption, as NielsenIQ data likely exclude some informal and online transactions, meaning flavored consumption frequencies for these age groups are likely larger.

[29] https://truthinitiative.org/research-resources/tobacco-industry-marketing/new-study-43-underage-e-cigarette-users-report.

Figure 4. Purchase Frequencies of Original and Flavored E-Cigarettes by Household Composition.

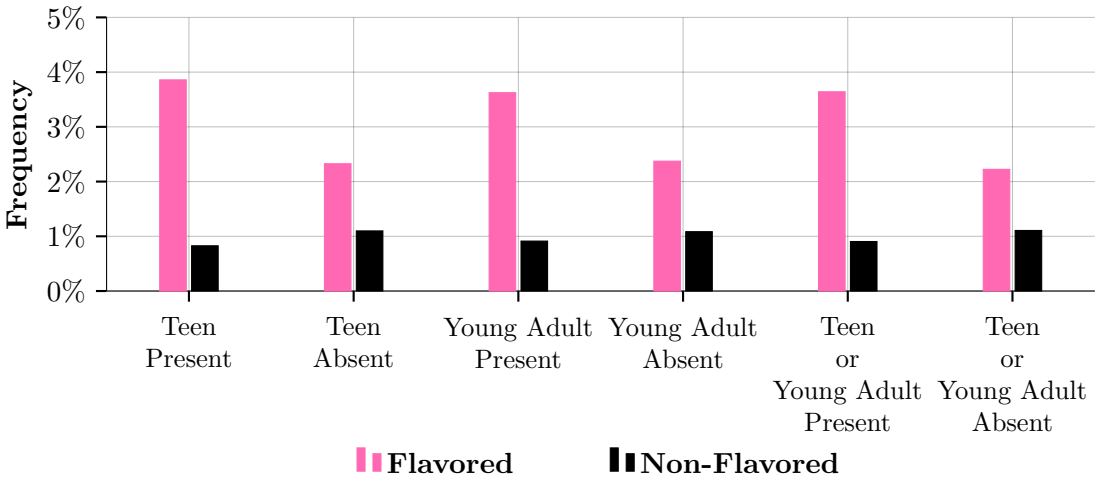

I further examine the relationship between household composition and the likelihood of purchasing flavored e-cigarettes in Table 2. The table reports three specifications of a linear probability model estimating whether a household purchases flavored e-cigarettes as a function of the presence of a teen or young adult, across two samples. Panel A covers all household-months in the sample. The most robust specification (column 3) shows that TYA households are 1.7 percentage points more likely to purchase flavored e-cigarettes in any given month, even unconditionally on having made an e-cigarette purchase at all. While small in absolute terms, this is a meaningful effect given that only about 2.5% of household-months involve a flavored e-cigarette purchase in the full sample. Panel B restricts to household-months in which an e-cigarette purchase was made, so the coefficient captures how household composition shapes the flavor choice conditional on buying e-cigarettes. In the most robust specification (column 3 of Panel B), households containing a teen or young adult are approximately 13.8 percentage points more likely to purchase flavored e-cigarettes. Given that 69.2% of all e-cigarette purchases in the sample are flavored, this implies that TYA households are about 20% more likely to purchase flavored e-cigarettes than the average e-cigarette-purchasing household.[30] Together, the two panels show that TYA households are both more likely to buy e-cigarettes at all and, conditional on doing so, more likely to choose flavored products.

[30] $(69.2 + 13.8)/69.2 = 83.0/69.2 \approx 1.20$, so TYA households purchase flavored e-cigarettes at a rate approximately 20% above the sample average conditional on an e-cigarette purchase.

Table 2. OLS Results for the Effect of a Teen or Young Adult in the Household on Purchasing Flavored E-Cigarettes.

Panel (A): All Household-Months for Inside Good Households

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Teen or Young Adult	0.014***	0.017***	0.017***
	(0.003)	(0.004)	(0.004)
HH Size		−0.001	−0.002
		(0.001)	(0.001)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	201,348	201,348	201,348

Panel (B): All Household Months Containing an E-Cig. Purchase

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Teen or Young Adult	0.138***	0.126***	0.138***
	(0.038)	(0.046)	(0.043)
HH Size		0.007	−0.001
		(0.015)	(0.014)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	7,207	7,207	7,207

Panel (A) includes all household-months for those households making at least one cigarette or e-cigarette purchase over the 2021 – 2023 time frame. Panel (B) restricts the sample to household-months containing e-cigarette purchases. Cluster-robust standard errors are at the household level and reported in parentheses. Significance levels: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3. Month-to-Month Category-Level Product Transitions

Panel (A): All Inside Good Households

Last Product Category Chosen	Current Product Category Chosen (%)					
	Outside Option	Cig. Only	Original E-Cig. Only	Flavored E-Cig. Only	Cig. & Original E-Cig.	Cig. & Flavored E-Cig.
Outside Option	**76.33**	21.78	0.42	1.23	0.06	0.18
Cig. Only	22.55	**76.67**	0.03	0.10	0.16	0.50
Original E-Cig. Only	25.48	0.93	**68.99**	1.80	2.73	0.06
Flavored E-Cig. Only	33.71	2.11	0.76	**60.29**	0.06	3.07
Cig. & Original E-Cig.	13.15	28.45	11.21	0.86	**41.38**	4.96
Cig. & Flavored E-Cig.	12.13	31.70	0.21	8.94	1.06	**45.96**

Panel (B): Inside Good Households with a Teen or Young Adult

Last Product Category Chosen	Current Product Category Chosen (%)					
	Outside Option	Cig. Only	Original E-Cig. Only	Flavored E-Cig. Only	Cig. & Original E-Cig.	Cig. & Flavored E-Cig.
Outside Option	**76.68**	20.90	0.49	1.67	0.03	0.23
Cig. Only	25.59	**73.20**	0.05	0.20	0.11	0.85
Original E-Cig. Only	36.24	1.74	**56.10**	4.18	1.74	0.00
Flavored E-Cig. Only	34.54	2.92	1.11	**57.70**	0.00	3.73
Cig. & Original E-Cig.	14.71	29.41	5.88	0.00	**44.12**	5.88
Cig. & Flavored E-Cig.	11.61	32.81	0.22	9.38	0.45	**45.54**

Rows denote the product category chosen in month $t-1$ and columns the category chosen in month t. Entries report transition probabilities in percent. The outside option corresponds to no purchase in a given month.

Table 3 presents month-to-month transition matrices of tobacco product choices, with Panel (a) covering all households that purchase tobacco-related products and Panel (b) covering those with a teen or young adult present. These transition structures offer insight into persistence and substitution dynamics across product categories. Both panels reveal a high frequency of non-purchasing months as well as strong within-category persistence, reflected in the large shares in the first column and along the main diagonal, respectively. Notably, this persistence is weaker among households with a teen or young adult, suggesting that younger individuals are more inclined to experiment with different products before settling on the one that sustains their consumption. Most interestingly, conditional on purchasing an e-cigarette in the previous month, the likelihood of switching to any type of cigarette in the current month is 2.45% for all households, but rises to

3.15% for those with a teen or young adult, a roughly 28.6% higher rate. This pattern aligns with the intuition that younger consumers are particularly susceptible to policies restricting flavored e-cigarettes, as such bans may unintentionally push them toward traditional cigarette use.[31]

I also examine the extent to which e-cigarette purchases involve brands that hold FDA marketing authorization. I classify a brand as FDA-authorized if it holds at least one FDA marketing order for any product.[32] This distinction is important for evaluating the practical reach of a flavor ban, since only FDA-authorized brands, which are currently limited to JUUL, Vuse, NJOY, and Logic, can be directly regulated through additional marketing orders. Figure 5 reports the share of flavored e-cigarette purchase months involving an FDA-authorized brand, separately for households with and without a teen or young adult present. Two patterns stand out. First, the FDA-authorized share has declined substantially over the sample period: from roughly 65% in 2021 to 42% in 2023 among households without a teen or young adult, and from 53% to 34% among those with one. Second, households with a teen or young adult consistently purchase FDA-authorized flavored e-cigarettes at lower rates. Together, these patterns suggest that a flavor ban limited to authorized brands would directly affect a shrinking and already minority share of flavored e-cigarette consumption, and would disproportionately miss the households policymakers are most concerned about. The trajectory in the data is consistent with national retail scanner evidence: by 2024, Truth Initiative estimates that unauthorized products account for approximately 86% of U.S. e-cigarette dollar sales, suggesting the unauthorized share has continued to grow beyond what the 2021–2023 HMS data capture.

This dominance reflects several structural advantages that unauthorized products hold over their authorized counterparts. FDA-authorized brands are restricted to tobacco- and menthol-flavored offerings, whereas unauthorized products are available in a wide array of fruit, candy, and dessert flavors that are particularly appealing to younger consumers. Unauthorized disposables also offer substantially greater e-liquid capacity — up to 16 milliliters or more per device, compared to roughly 1.8 milliliters for authorized pod systems — and operate at higher wattage with adjustable power settings, delivering

[31] Chen and Reinhardt (2025) produce a similar matrix to that of Panel (a) in Table 3 using weekly level purchases of the NielsenIQ Household Scanner over 2015 – 2019, inclusive. They find much higher transition probabilities of e-cigarettes into combustible cigarettes. Their finding may be due to e-cigarettes being a newer product at the time, relative to the time period my data spans, and tobacco product users still experimenting with e-cigarettes before largely transitioning from traditional cigarettes.

[32] Throughout this paper, I classify a brand as FDA-authorized if it holds at least one Premarket Tobacco Product Application (PMTA) FDA marketing order for any product in its portfolio. As of 2023, four brands meet this criterion: JUUL, Vuse, NJOY, and Logic. Importantly, this is a brand-level classification, not a product-level one. The FDA has granted marketing orders only for tobacco- and menthol-flavored products from these brands; their flavored offerings (e.g., fruit, candy, or dessert varieties) do not themselves hold marketing orders. The brand-level classification is appropriate for the counterfactual exercises in this paper because firms with existing marketing orders face regulatory exposure that creates strong compliance incentives across their entire product line: non-compliance on unauthorized products risks jeopardizing their existing authorizations. Brands without any marketing orders face no such leverage, which is why enforcement against them has been largely ineffective.

more vapor and nicotine per device at a lower effective cost per puff. Together, these product-level differences make unauthorized e-cigarettes meaningfully more attractive to consumers along dimensions beyond flavor alone, which helps explain why enforcement efforts targeting authorized brands have had limited impact on overall consumption.

Figure 5. Share of Flavored E-Cigarette Purchase Months by FDA Authorization Status and Teen or Young Adult Presence.

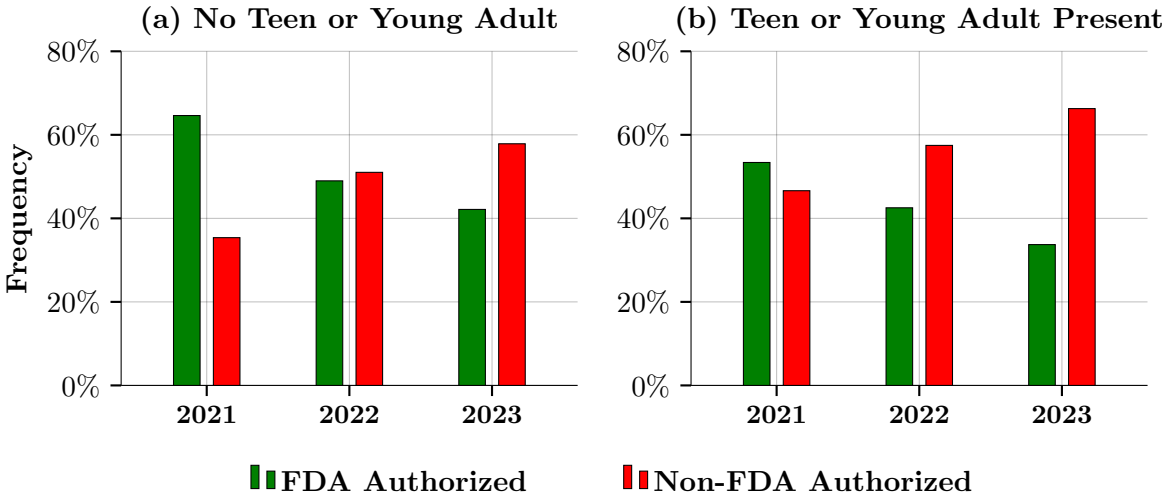

Finally, to further assess the dynamics of tobacco consumption, I conduct an abstention spell analysis that tracks households through abstention spells, defined as months in which a household purchased a tobacco product in the previous month but not the current one, and measures the number of months until the household next purchases. Table A.6 reports these statistics separately for cigarettes and e-cigarettes, and Figure 6 plots the empirical cumulative distribution of months to resumption among spells that ended in a return to purchasing. Two features stand out. First, return to purchasing is pervasive: 72.5% of cigarette abstention spells and 59.8% of e-cigarette abstention spells end in a return to purchasing within the panel window. Second, return to purchasing is fast: among those who return, roughly half of cigarette abstainers do so within a single month (49.9%), as do 45.2% of e-cigarette abstainers, and roughly three-quarters of both groups return within three months. The front-loaded shape of this distribution is difficult to reconcile with standard exponential discounting, under which a household that finds it optimal to stop purchasing should remain stopped absent a sufficiently large shock. By contrast, this pattern arises naturally under present-biased preferences, where a household may plan to abstain but re-optimizes when the future becomes the present.[33]

[33]These abstention spells are defined at the purchase level, not as self-reported cessation attempts. One concern is that some spells may reflect temporary interruptions such as stockpiling or travel rather than genuine cessation attempts. However, the monthly aggregation of the data substantially mitigates this: a household would need to stockpile an entire month's supply of tobacco or travel for a full month without purchasing, both of which are uncommon for regular consumers of addictive products.

Figure 6. Empirical Cumulative Distribution of Months to Resumption, Among Abstention Spells That Ended in a Return to Purchasing.

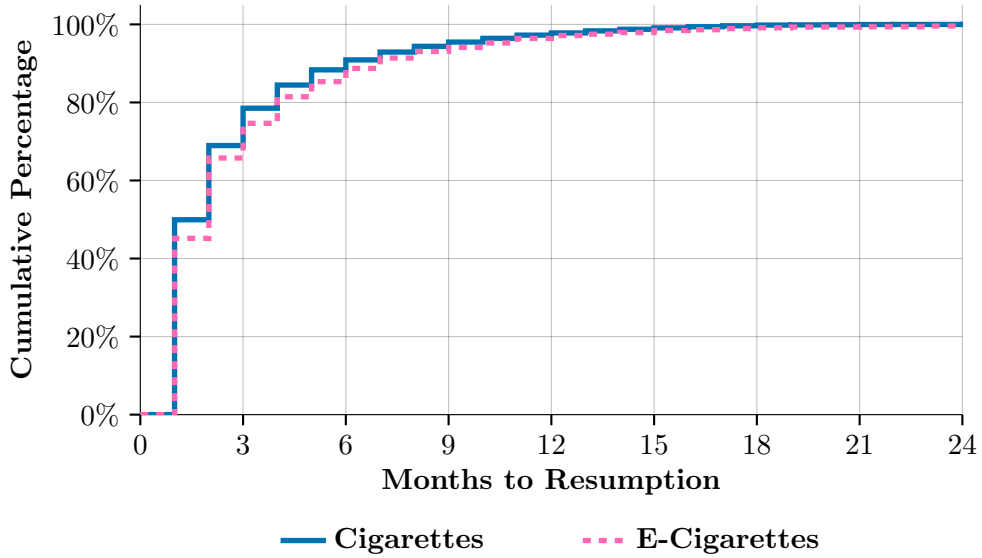

Taken together, the reduced-form evidence presented above motivates both the dynamic structure of the model and several of its specific features developed in Section V. When current consumption raises future addiction stocks, which in turn raise the cost of abstention, the demand effects of a flavor policy compound over time in ways a static model model cannot recover. A dynamic model is necessary for tracing out the long-run addiction and welfare consequences of the policies studied here. The disproportionate purchasing of flavored e-cigarettes by households with teens and young adults (Table 2) motivates heterogeneous flavor preferences that vary with household composition. The strong within-category persistence revealed in the transition matrices (Table 3), combined with the non-trivial cross-category substitution (and in particular the higher e-cigarette-to-cigarette transition rate among households with younger members), motivates both the addiction stock that links current consumption to future utility and the multi-product choice set that allows for substitution across cigarettes, e-cigarettes, and their bundles. The declining share of FDA-authorized brands (Figure 5) motivates the distinction between authorized and unauthorized flavored alternatives in the counterfactual analysis. Finally, the abstention spell analysis (Figure 6) motivates the quasi-hyperbolic discounting specification used in the welfare analysis: the sharp front-loading of returns to purchasing is inconsistent with a fully rational consumer who finds abstention optimal and sticks to that plan, but arises naturally when consumers are present-biased and re-optimize once the future becomes the present.

Furthermore, I perform a reduced-form test for stockpiling in Table A.5, which suggest this is not a prevalent feature of data.

V. Structural Model

V.1. *Overview*

The rational addiction framework models consumption choices over time using a dynamic optimization problem. Consumers maximize the expected discounted sum of flow utilities, where current-period utility depends on consumption, nicotine intake, evolving nicotine addiction and flavored habit stocks, expenditures, and product and household attributes. The decision-making process within each period follows the timeline in Figure 7.

Figure 7. Within Period Timing for Consumers

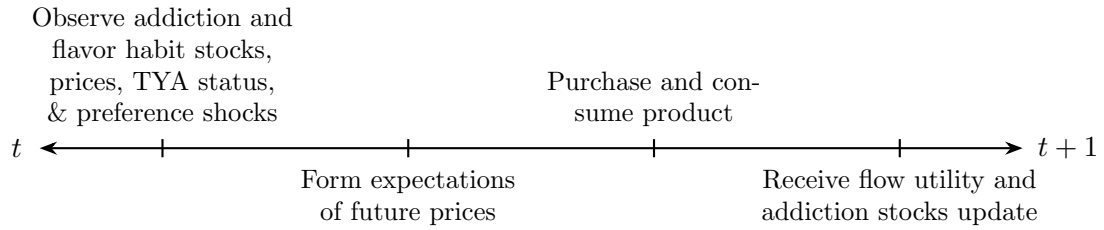

At the start of each period t, the consumer observes the current state, which consists of: (1) their nicotine addiction stocks, (2) their flavored habit stock, (3) the prevailing prices across product categories, (4) whether a teen or young adult is in the household, and (5) a vector of idiosyncratic preference shocks across alternatives. Nicotine addiction and the flavored habit stock both decay over time unless reinforced by new consumption. Preference shocks capture unobservable factors such as stress, social influences, and situational cravings that shift preferences in ways I cannot observe. The consumer then forms expectations about future prices and teen-or-young-adult transitions. Based on this information, they choose an alternative to purchase and consume within the period. Finally, the consumer receives flow utility from their decision.

V.2. *Choice Set*

The choice set comprises two nested parts: product categories and specific alternatives within each category. By organizing choices into categories, this simplifies notation and significantly reduces the computational burden due to the state space. The latter is further discussed in Section X.4.2..

There are four product categories: the outside option ($\mathcal{O}$, no purchase), combustible cigarettes ($\mathcal{C}$), e-cigarettes ($\mathcal{E}$), and bundles of cigarettes and e-cigarettes purchased together ($\mathcal{B}$). Within the e-cigarette category, products are further divided into three subcategories: original (non-flavored) e-cigarettes ($\mathcal{E}_o$), flavored e-cigarettes sold by brands that have not received FDA marketing authorization ($\mathcal{E}_{f,na}$), and flavored e-cigarettes sold by brands that received FDA marketing authorization ($\mathcal{E}_{f,a}$).

Nested within each category are specific alternatives differentiated by purchasable quantity. Cigarettes are offered in twelve quantity bins; each of the three e-cigarette subcategories has seven quantity bins; and bundles pair cigarettes with e-cigarettes of each subcategory at two intensity levels, yielding six bundle alternatives.[34] I include bundles because roughly 26.6% of household-months containing e-cigarette purchases also contain cigarette purchases, and modeling joint purchases allows for complementary consumption in the spirit of Gentzkow (2007).[35]

The quantity bins for each category are chosen based on the empirical purchase frequency distributions shown in Figure 8. I use twelve bins for cigarettes: 1, 2, 3–4, 5–9, 10, 11–19, 20, 21–29, 30, 31–39, 40, and 41+ packs; and seven bins for each of the three e-cigarette subcategories (original, non-FDA authorized flavored, and FDA authorized flavored): 0–5, 5–10, 10–15, 15–20, 20–30, 30–50, and 50+ mL. Together with the outside option and six bundle alternatives, this yields 40 total alternatives.

[34] While cigarette-only and e-cigarette-only alternatives are differentiated across several quantity bins, bundles use a coarser disaggregation consisting of two intensity levels crossed with three e-cigarette subcategories, which reflects data constraints. Joint purchases of cigarettes and e-cigarettes comprise approximately 0.9% of household-months. Disaggregating bundles further by e-cigarette quantity would yield extremely small shares for those cells. Since the primary counterfactual of a ban on flavored e-cigarettes eliminates the flavored bundle alternatives entirely regardless of quantities involved, the relevant margin is where flavored bundle purchasers substitute when those options are removed, not how they adjust quantities within bundles.

[35] This nested category–alternative structure offers several advantages. First, it parallels the inclusive value sufficiency (IVS) logic of Gowrisankaran and Rysman (2020) by grouping products into economically meaningful categories and sub-categories, thereby collapsing a potentially huge alternative-level price state space into a tractable set of category-level state variables. This not only simplifies notation, but also reduces the computational complexity of solving the dynamic programming problem, as future states can be summarized by a small set of categories rather than the full price vector for every alternative that is differentiated not only by potential flavor, but also purchasable quantity. Second, it aligns naturally with consumer substitution patterns observed in nicotine product markets, where cross-substitution is typically strongest within categories (e.g., across original, non-FDA authorized flavored, and FDA authorized flavored e-cigarettes) and weaker across them as illustrated in Table 3. Unlike the constant continuation value assumption often paired with IVS, however, this specification does not force all products within a category to share the same continuation value. Quantity choices affect the evolution of the addiction stock, which in turn alters future utility, so even products in the same category can have different continuation values conditional on the state. This preserves important within-category dynamics while still delivering large computational savings relative to allowing the state space to depend on particular alternatives. In standard IVS and CCV setups, continuation values are identical for all products within a category, which can understate within-category substitution dynamics and mute the effect of policies that change state transitions (e.g., addiction evolution). By relaxing CCV within categories, the model captures richer behavioral responses to such policies.

Figure 8. Distributions of Monthly Packs of Cigarettes Purchased (Panel (A)) and Monthly Milliliters of E-Liquid Purchased (Panel (B)).

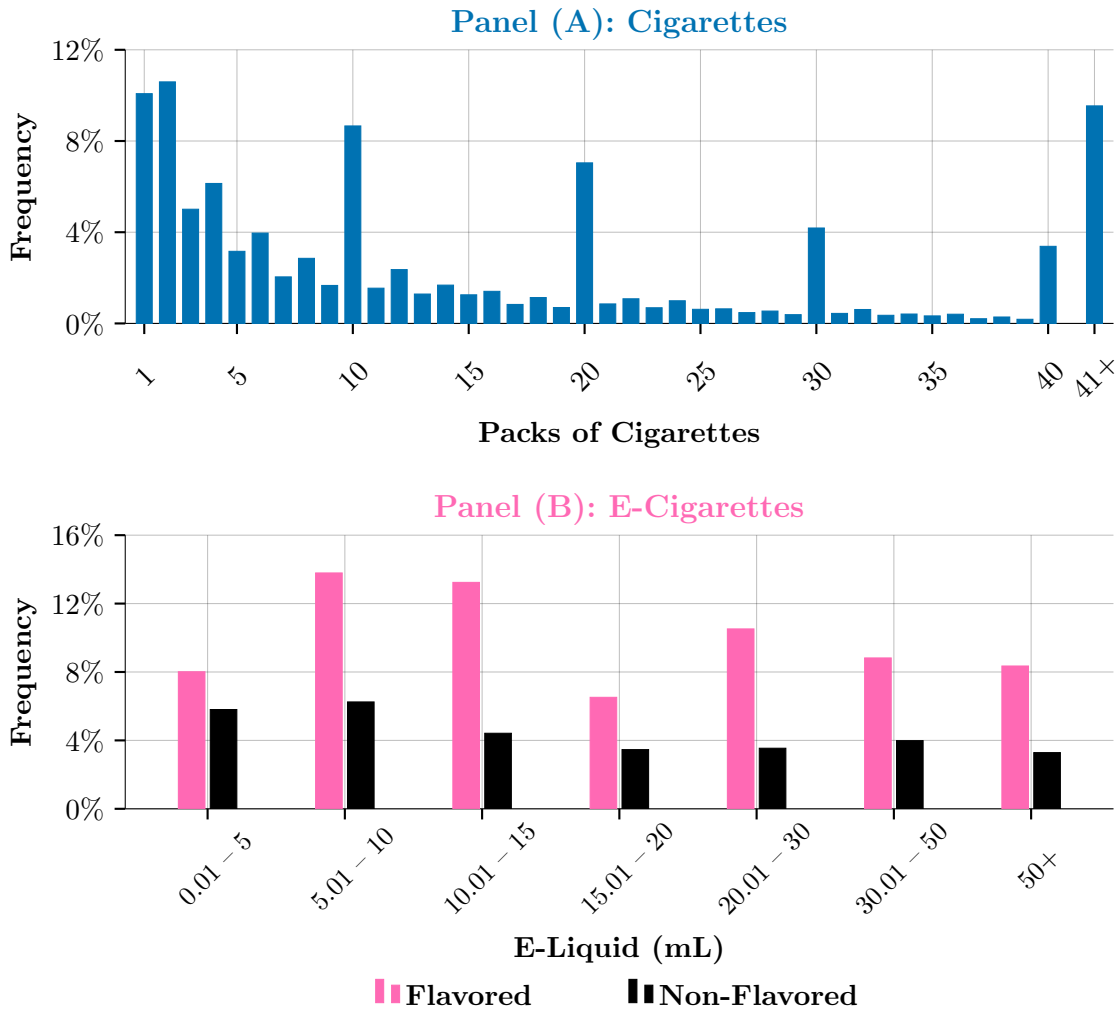

The final bar for packs of cigarettes purchased reflects the aggregate share of 41 or more packs purchased. Similarly, the bar at 50+ for milliliters of e-liquid purchased reflects the aggregate share of over 50 milliliters of e-liquid purchased.

V.3. *Purchase Decisions and Consumption*

In each period t, consumer i observes the state $\boldsymbol{x}_{it}$ and a vector of idiosyncratic preference shocks $\boldsymbol{\varepsilon}_{it} = (\varepsilon_{i1t}, \ldots, \varepsilon_{i,|\mathcal{J}|,t})$ across all alternatives, forms expectations of future prices, and chooses an alternative $j \in \mathcal{J}$. This choice is represented by $\boldsymbol{d}_{it} = (d_{i1t}, d_{i2t}, \ldots, d_{i,|\mathcal{J}|,t})$, where

$$d_{ijt} = \begin{cases} 1, & \text{if consumer } i \text{ chooses alternative } j \text{ at time } t, \\ 0, & \text{otherwise.} \end{cases}$$

I do not model inventory accumulation or stock-outs, consistent with the policy

evaluation focus of this paper.[36] Under this assumption, all purchased quantity is consumed within the period. Each alternative $j \in \mathcal{J}$ encodes a fixed cigarette quantity q_{iCt} and e-cigarette quantity q_{iEt} when consumed by household i in period t, where for single-category alternatives one of these is zero and for bundles both are positive. Each alternative is also associated with a level of absorbed nicotine intake n_{ijt}, described in Section V.4..

V.4. *Addiction Stock and Nicotine Intake*

The flow utility depends not only on current consumption and prices but also on the consumer's accumulated addiction stock. Modeling addiction as a stock that accumulates in nicotine units is what allows the model to trace out long-run addiction trajectories.

The standard rational addiction model of Becker and Murphy (1988) represents addiction with a single stock governed by one decay parameter. However, a single decay rate cannot simultaneously capture two well-established features of nicotine addiction: (i) acute withdrawal symptoms that peak within the first week and largely resolve within one to two months (Mamede et al. (2007); Cosgrove et al. (2009); McLaughlin, Dani, and De Biasi (2015)), and (ii) elevated relapse risk that persists for months or years after physiological withdrawal has resolved, driven by learned associations between nicotine reward and routines such as meals, work breaks, alcohol consumption, and emotional stress (Hughes, Keely, and Naud (2004); Hughes, Peters, and Naud (2008); Krall, Garvey, and Garcia (2002)). A fast single stock would correctly predict that withdrawal fades quickly but would understate the long-run relapse hazard; a slow single stock would capture persistent relapse risk but would overstate withdrawal duration. To resolve this tension, I decompose addiction into two components: a fast stock a^f_{it} that captures short-run physiological craving, and a slow stock a^s_{it} that captures long-run behavioral dependence. Each stock evolves according to its own law of motion:

$$a^s_{i,t+1} = (1-\psi_1)\cdot a^s_{it} + n_{iCt} + n_{iEt},$$
$$a^f_{i,t+1} = (1-\psi_2)\cdot a^f_{it} + n_{iCt} + n_{iEt},$$

where $\psi_1 = 0.10$ and $\psi_2 = 0.90$ are the decay rates for the slow and fast stocks, respectively, and n_{iCt} and n_{iEt} are the absorbed nicotine from the cigarette and e-cigarette components of the chosen alternative.[37]

The addiction level entering the flow utility is the unweighted average of the two

[36] While inventory dynamics are sometimes modeled in the tobacco demand literature (Hendel and Nevo (2006); Gordon and Sun (2015); Chen and Rao (2020)), including them would substantially increase the state space without directly informing the flavor ban counterfactuals.

[37] Absorption rates differ across product categories. Moreover, all cigarettes typically contain the same amount of nicotine irrespective of brand. For further information, see Section X.7..

stocks:

$$a_{it} = \frac{a_{it}^f + a_{it}^s}{2}.$$

The fast stock, with $\psi_2 = 0.90$, retains only 10% of its value after one month and is effectively zero after two months, aligning with the timeline of acute withdrawal. The slow stock, with $\psi_1 = 0.10$, retains 53% of its value at six months and 28% at one year, aligning with the gradual decay of relapse hazard rates among long-term abstainers. Both values are calibrated from the biomedical literature on nicotine addiction rather than estimated jointly with the structural parameters; a detailed discussion is provided in Section X.6..[38]

The standard rational addiction specification uses raw consumption quantities in the law of motion, effectively treating all consumption units as homogeneous in their addictive impact. But addiction is physiologically driven by absorbed nicotine, not consumption per se. Since a pack of cigarettes and a milliliter of e-liquid deliver very different amounts of nicotine, and the data allow me to infer nicotine strength for each product, incorporating n_{iCt} and n_{iEt} provides a more precise law of motion for the addiction stock. The conversion from nicotine concentration to absorbed nicotine is detailed in Section X.7..

V.5. *Flavored Habit Stock*

The nicotine addiction stocks capture physiological dependence that is common across all nicotine-delivering products. However, the descriptive evidence in Section IV. reveals strong within-category persistence in flavored e-cigarette purchasing that cannot be fully explained by nicotine addiction alone: households that purchase flavored e-cigarettes in one month are substantially more likely to purchase them again the next month, even conditional on their overall nicotine consumption level. To capture this flavor-specific state dependence, I introduce a separate flavored habit stock a_{it}^{flav} that tracks the household's accumulated history of flavored purchasing independently of its nicotine addiction state.

[38] Monte Carlo simulations revealed that ψ_1 and ψ_2 are poorly identified when estimated jointly with the structural parameters. The problem is intuitive: in scanner data, the only information about addiction comes from watching what households buy over time. A household that keeps smoking month after month could be explained by many combinations of a fast-decaying stock (ψ_2 large) and a slow-decaying stock (ψ_1 small), since both the short-run craving dynamics and the long-run dependence trajectory are determined by these two decay rates jointly. When ψ_1 and ψ_2 are free parameters, the optimizer can rationalize the same observed purchasing persistence through a range of fast-slow combinations, leaving the individual rates poorly pinned down. To separately identify both decay rates, one would need to observe how consumption rebounds after disruptions occurring at different timescales, such as short hospital stays that interrupt nicotine intake (revealing ψ_2) versus long-term cessation programs that reveal the slow persistence of relapse risk (revealing ψ_1). Since scanner data contain no such observable exogenous disruptions, I calibrate ψ_1 and ψ_2 from the biomedical literature instead.

The flavored habit stock evolves according to:

$$a_{i,t+1}^{\text{flav}} = (1 - \psi_3) \cdot a_{it}^{\text{flav}} + \mathbb{1}\left\{j_{ilt} \in \mathcal{E}_f\right\},$$

where $\psi_3 \in (0, 1)$ is the decay rate, estimated jointly with the other structural parameters, and $\mathbb{1}\{j_t \in \mathcal{E}_f\}$ is an indicator for whether the household chose a flavored e-cigarette alternative in period t. Unlike the nicotine addiction stocks, which accumulate in milligrams of absorbed nicotine and are therefore dose-dependent, the flavored habit stock tracks the frequency rather than the intensity of flavored purchasing. The nicotine addiction stocks already capture the intensity dimension of flavored e-cigarette purchasing, since any flavored product consumed contributes nicotine to the addiction stock in proportion to the quantity purchased. What the addiction stocks cannot capture is the categorical pattern of choosing flavored products specifically. A household that buys a large bottle of flavored e-liquid builds a larger nicotine addiction stock than one that buys a small pod, but both contribute equally to the flavored habit stock because both made a flavored purchase. The flavor-specific inertia is reinforced by the act of choosing flavored, not by how much was consumed.

This specification is important for the flavor ban counterfactual because it provides a mechanism for flavor-specific persistence that is distinct from nicotine addiction. Without the flavored habit stock, the model would attribute all within-category persistence to nicotine dependence, which is common across products and therefore cannot explain why flavored e-cigarette users are more likely to continue purchasing flavored products specifically. By separating flavor-specific habits from nicotine addiction, the model can generate more realistic substitution patterns when flavored products are removed: the flavored habit stock determines the initial reluctance to switch, while the nicotine addiction stocks determine whether and how quickly consumers substitute toward other nicotine products rather than exiting the market entirely.

V.6. *Price Evolution*

Following Erdem, Imai, and Keane (2003), I assume consumers believe prices evolve according to a first-order autoregressive process:

$$p_{C,t+1} = \phi_{0C} + \phi_{1C} \cdot p_{Ct} + \eta_{C,t+1},$$
$$p_{E,t+1} = \phi_{0E} + \phi_{1E} \cdot p_{Et} + \eta_{E,t+1},$$

with

$$\begin{bmatrix} \eta_{C,t+1} \\ \eta_{E,t+1} \end{bmatrix} \sim \mathbb{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma \right),$$

where p_{kt} denotes the per-unit price for category $k \in \mathcal{K}$ at time t, ϕ_{0k} is a category-specific intercept, ϕ_{1k} captures price persistence, and $\eta_{k,t+1}$ is a mean-zero normally distributed price shock. This process operates at the category level: all alternatives $j \in \mathcal{J}_k$ share the same price parameters.[39] The AR(1) is estimated separately for each category, so price expectations are category-specific.

The covariance matrix Σ allows for correlated price shocks across categories, reflecting the fact that cigarette and e-cigarette prices share common drivers such as retail margins, state tax environments, and broader macroeconomic conditions affecting consumer goods. Allowing for this correlation ensures that consumers' price expectations correctly account for the tendency of prices across categories to move together, rather than treating them as independent. This is distinct from the equilibrium price responses to demand shifts induced by a policy, which are handled by the supply-side model.

V.7. *Flow Utility*

Having defined the choice set, state variables, and their laws of motion, I now specify how these components enter the consumer's period payoff. In each period t, consumer i observes state $\boldsymbol{s}_{it} = (\boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}) = \left(h_{it}, a^f_{it}, a^s_{it}, a^{\text{flav}}_{it}, \boldsymbol{p}_{it}, \boldsymbol{\varepsilon}_{it}\right)$, where $\boldsymbol{x}_{it} = \left(h_{it}, a^f_{it}, a^s_{it}, a^{\text{flav}}_{it}, \boldsymbol{p}_{it}\right)$ is the observed portion of the state consisting of the teen-or-young-adult indicator $h_{it} \in \{0, 1\}$, the fast and slow addiction stocks a^f_{it} and a^s_{it}, the flavored habit stock a^{flav}_{it}, and the category-level price vector $\boldsymbol{p}_{it} = (p_{iCt}, p_{iEt})$. The average addiction stock entering the flow utility is $a_{it} = (a^f_{it} + a^s_{it})/2$. The vector $\boldsymbol{\varepsilon}_{it}$ collects unobserved idiosyncratic shocks to utility across all alternatives. The flow utility from choosing alternative $j \in \mathcal{J}$ for a consumer of latent type $l \in \{1, 2\}$ is

$$
\begin{aligned}
u_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right) = {} & \alpha_C \cdot q_{iCt} + \alpha_E \cdot q_{iEt} + \alpha_{CE} \cdot q_{iCt} \cdot q_{iEt} \\
& + \left(\lambda_1 + \lambda_2 \cdot \mathbb{1}\left\{h_{it} = 1\right\}\right) \cdot \mathbb{1}\left\{j \in \mathcal{E}_f\right\} \\
& + \left(\lambda_3 + \lambda_4 \cdot \mathbb{1}\left\{h_{it} = 1\right\}\right) \cdot \mathbb{1}\left\{j \in \mathcal{E}_{f,a}\right\} \\
& + \gamma_1 \cdot a_{it} \cdot \mathbb{1}\left\{j = \mathcal{O}\right\} \\
& + \gamma_2 \cdot a^{\text{flav}}_{it} \cdot \mathbb{1}\left\{j \in \mathcal{E}_o\right\} \\
& + \gamma_3 \cdot a^{\text{flav}}_{it} \cdot \mathbb{1}\left\{j \in \mathcal{C}\right\} \\
& + \gamma_4 \cdot a^{\text{flav}}_{it} \cdot \mathbb{1}\left\{j = \mathcal{O}\right\} \\
& + \omega_C \cdot p_{iCt} \cdot q_{iCt} + \omega_E \cdot p_{iEt} \cdot q_{iEt} \\
& + \xi_{lk(j)} + \varepsilon_{iljt},
\end{aligned}
$$

where $\boldsymbol{\theta}^u$ is the vector of structural utility parameters to be estimated. When consumer i does not purchase a product at time t (i.e., chooses the outside option $j = \mathcal{O}$), all

[39] Category-level price expectations substantially reduce the state space. Since alternatives in $\mathcal{J}$ are differentiated by quantity, tracking price expectations for each product-quantity pair would make the dynamic program computationally intractable. See Section X.4.2. for further discussion.

consumption, price, flavor, and fixed effect terms are zero except the addiction and flavored withdrawal costs, so the flow utility reduces to

$$u_{il0t} = \gamma_1 \cdot a_{it} + \gamma_4 \cdot a_{it}^{\text{flav}} + \varepsilon_{i0t}.$$

The consumption parameters α_C, α_E, and α_{CE} capture the direct utility from tobacco product use. The terms $\alpha_C \cdot q_{iCt}$ and $\alpha_E \cdot q_{iEt}$ represent the baseline marginal utilities of cigarette and e-cigarette consumption, respectively, where I expect $\alpha_C > 0$ and $\alpha_E > 0$. The interaction $\alpha_{CE} \cdot q_{iCt} \cdot q_{iEt}$, which is nonzero only for bundle alternatives where both quantities are positive, captures whether the two products are complements ($\alpha_{CE} > 0$) or substitutes ($\alpha_{CE} < 0$) when consumed within the same period. The price terms $\omega_C \cdot p_{iCt} \cdot q_{iCt}$ and $\omega_E \cdot p_{iEt} \cdot q_{iEt}$ enter as expenditure-style disutility: the price cost scales with the quantity purchased, so larger purchases carry greater expenditure disutility. I expect $\omega_C < 0$ and $\omega_E < 0$. Separating the price coefficients allows the model to capture differential price responsiveness across product categories. Cigarette alternatives receive only the cigarette expenditure term, e-cigarette alternatives receive only the e-cigarette term, and bundles receive both.

The flavor parameters λ_1 through λ_4 capture the utility premium associated with flavored e-cigarette products and how it varies with household composition. The baseline effect λ_1 applies to all flavored alternatives $j \in \mathcal{E}_f$, while λ_2 captures the differential effect for households with a teen or young adult member ($h_{it} = 1$). The total flavor effect is λ_1 for non-TYA households and $\lambda_1 + \lambda_2$ for TYA households. The parameters λ_3 and λ_4 provide an analogue for FDA-authorized flavored alternatives $j \in \mathcal{E}_{f,a}$: λ_3 captures the baseline utility difference between FDA-authorized and non-FDA flavored products, and λ_4 captures whether this difference varies for TYA households. Together, these four parameters determine the immediate utility loss when flavored alternatives are removed from the choice set and thus govern the direction and magnitude of consumer substitution following a flavor ban.

The parameter γ_1 governs the dynamic interaction between the accumulated addiction stock a_{it} and the decision to abstain from nicotine, embedding the model in the rational addiction framework of Becker and Murphy (1988) and the habit formation tradition of Stigler and Becker (1977). The abstention cost term $\gamma_1 \cdot a_{it} \cdot \mathbb{1}\{j = \mathcal{O}\}$ captures the disutility of not consuming nicotine when addicted. This represents the full cost of abstention (combining both withdrawal pain and the forgone reinforcement from continued use) that makes cessation increasingly difficult as the stock accumulates, similar in spirit to the stock-dependent utility from previously consumed durables in Gowrisankaran and Rysman (2012). I expect $\gamma_1 < 0$. The self-reinforcing cycle characteristic of addictive behavior operates through this single channel: consumption builds the addiction stock, which raises the cost of abstention, which makes continued purchasing more likely, which further builds the stock.

The parameters γ_2, γ_3, and γ_4 introduce state dependence through interactions with the flavored habit stock a_{it}^{flav}, and all three are expected to be negative. The term $\gamma_2 \cdot a_{it}^{\text{flav}} \cdot \mathbb{1}\{j \in \mathcal{E}_o\}$ penalizes choosing original (non-flavored) e-cigarette and original bundle alternatives for households with an elevated flavored habit stock, capturing the reluctance of habitual flavored purchasers to switch to unflavored products within the e-cigarette category. The term $\gamma_3 \cdot a_{it}^{\text{flav}} \cdot \mathbb{1}\{j \in \mathcal{C}\}$ imposes an analogous penalty on choosing cigarette and cigarette-containing bundle alternatives. The term $\gamma_4 \cdot a_{it}^{\text{flav}} \cdot \mathbb{1}\{j = \mathcal{O}\}$ penalizes market exit for households with a high flavored habit stock, capturing a flavored withdrawal cost that makes cessation more difficult the more entrenched the flavored purchasing habit. Putting all four together, γ_1 governs the cost of abstaining from nicotine given addiction; γ_2 governs the attractiveness of unflavored e-cigarettes relative to flavored alternatives; γ_3 governs the attractiveness of cigarettes relative to flavored e-cigarettes; and γ_4 governs the cost of exiting the nicotine market entirely given a high flavored habit stock. When flavored alternatives are removed, households face a choice between unflavored e-cigarettes (penalized by γ_2), cigarettes (penalized by γ_3), and market exit (penalized jointly by γ_1 through nicotine addiction and γ_4 through the flavored habit stock). The relative magnitudes of these penalties shape whether post-ban substitution flows toward unflavored e-cigarettes, toward cigarettes, or out of the nicotine market entirely.

To capture permanent unobserved heterogeneity in tobacco propensity, the flow utility allows for $L = 2$ latent consumer types, $l \in \{1, 2\}$, each with its own category-specific intercepts while sharing all other structural parameters. The type-specific intercepts $\xi_{lk(j)}$ capture baseline utility differences across product categories that the consumption, flavor, and addiction terms do not explain. The category index $k(j)$ maps each alternative to cigarettes (C), e-cigarettes (E), or bundles (CE), yielding six intercepts in total ($\xi_{1C}, \xi_{1E}, \xi_{1CE}, \xi_{2C}, \xi_{2E}, \xi_{2CE}$). Two consumers may respond identically to prices, flavors, and addiction stocks yet differ substantially in their baseline propensity to choose cigarettes over e-cigarettes, reflecting unobserved differences in category-level preferences.[40] The treatment of type membership as latent, the mixing weight specification, and the distributional assumption on ε_{iljt} are discussed in Section VI.4..

VI. Identification and Estimation

VI.1. *Identification*

I now discuss how each group of structural parameters is identified, emphasizing the variation in the data that pins down each object and the assumptions required.

Consumption parameters. The consumption parameters α_C, α_E, and α_{CE} are identified from the relative frequency with which households choose higher- versus lower-

[40] Without type-varying intercepts, the model would force all consumers to share the same baseline category valuation, potentially attributing heterogeneity in category choice to price sensitivity or addiction parameters and biasing counterfactual predictions.

quantity alternatives within the same product category. Consider two cigarette alternatives that differ only in quantity, such as 5 packs versus 20 packs in a given month. A household that regularly selects the 20-pack alternative over the 5-pack, holding prices fixed, reveals a higher marginal utility of cigarette consumption, and this variation identifies α_C. The same logic applies to α_E across e-cigarette volume bins. The bundle interaction α_{CE} is identified from how frequently households choose bundle alternatives relative to single-category alternatives. If consuming both products together delivers less utility than the sum of each consumed separately, households will avoid bundles even when both products are available, revealing $\alpha_{CE} < 0$ and indicating that cigarettes and e-cigarettes are within-period substitutes.

Two endogeneity concerns arise. Along the extensive margin, households choose which product category to purchase, and unobserved preferences for cigarettes versus e-cigarettes could bias the consumption parameters if not accounted for. The type-specific category fixed effects $\xi_{lk(j)}$ address this directly: they absorb all time-invariant unobserved utility differences across product categories for each latent type, so that α_C and α_E are identified purely from within-category quantity variation rather than from cross-category sorting driven by unobserved tastes. Without these fixed effects, a segment of consumers who prefer e-cigarettes for reasons unrelated to nicotine content, such as health perceptions or social norms, would spuriously inflate $\widehat{\alpha}_E$. Along the intensive margin, households decide how much to buy conditional on category, and addiction complicates this: a highly addicted household buys more not just because consumption is intrinsically valued but because the addiction stock raises the disutility of purchasing less. The explicit modeling of the nicotine addiction stock as a state variable at least partially accounts for this by separating the addiction-driven component of quantity demand from the consumption utility parameters.

Price sensitivity. Identification of the expenditure coefficient ω relies on variation in retail prices across households and over time. A natural concern is price endogeneity: if firms set higher prices for products with unobservably high demand, the estimated ω would be biased toward zero. In standard BLP-style settings (Berry, Levinsohn, and Pakes (1995)), this arises because the unobserved product characteristic enters both utility and the firm's pricing decision. In my setting, this concern is mitigated by three features. First, the category fixed effects $\xi_{lk(j)}$ absorb all time-invariant unobserved demand differences across categories and types. This means that any remaining variation in utility across alternatives within a category comes from quantity differences, not from unobserved product characteristics that firms could price against.[41] Second, the price a household faces on the shelf is not a response to that household's demand. Philip

[41] This argument would break down if the model included product-level differentiation within a category beyond quantity. For example, if individual brands or UPCs entered utility with their own unobserved characteristics. In that case, a brand with high unobserved quality could command a price premium. Addressing this would require instrumental variables, such as BLP-style cost shifters or Hausman-type instruments. Since my alternatives within categories are defined by quantity rather than individual products, this concern does not arise in the current specification.

Figure 9. Share of Flavored and Unflavored E-Cigarette Purchases.

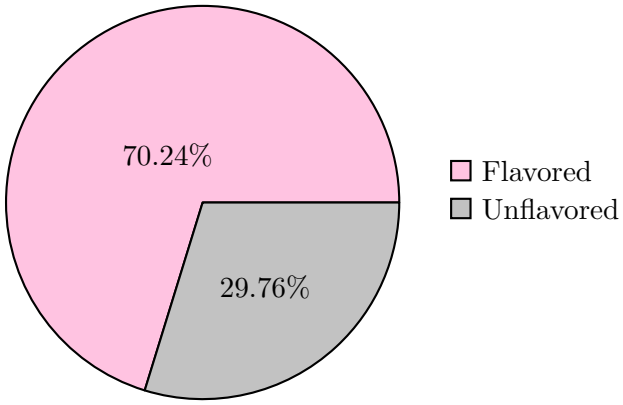

Morris does not observe that a particular household in Tucson had a stressful week and raise the price of Marlboros at their local gas station. Wholesale prices are set nationally or regionally by manufacturers, and retail markups reflect store-level costs, state excise taxes, and margins, none of which respond to individual household demand shocks. Third, the dominant source of cross-state price variation is state excise tax policy, which legislatures set for fiscal and public health reasons rather than in response to contemporaneous demand. Together, these features imply that the price variation identifying ω is driven primarily by cost-side factors rather than demand shocks.

Flavor parameters. The flavor parameters λ_1 through λ_4 are identified from how households sort across flavored, unflavored, and FDA-authorized e-cigarette alternatives. The baseline flavor premium λ_1 comes from the overall split between flavored and unflavored e-cigarette purchases: roughly 70% of e-cigarette purchases are flavored (Figure 9), and this share pins down λ_1 once prices, addiction, and type fixed effects are accounted for. The TYA interaction λ_2 is identified from the gap in flavored shares between households with and without teens or young adults: TYA households buy flavored products at a higher rate even after controlling for prices and addiction (Figure 4), and λ_2 captures the portion of this gap attributable to a genuine flavored product premium for TYA households. The FDA parameters λ_3 and λ_4 come from within-flavored variation: among households that buy flavored e-cigarettes, the split between FDA-authorized brands and non-FDA brands, and how that split varies with TYA status (Figure 5), pins down these two coefficients.

Nicotine addiction parameter. The abstention cost γ_1 is identified from how the probability of choosing the outside option varies across households with different nicotine addiction stocks. A household that has been buying tobacco products every month for a year has accumulated a large addiction stock; it should be much less likely to go without nicotine in a given month than a household that buys only occasionally. The variation in abstention rates across households with different purchase histories pins down γ_1. Since the nicotine addiction stock enters the outside option utility only through γ_1 (all inside

alternatives carry zero nicotine abstention cost), identification of the addiction effect on market exit comes from the inside-versus-outside margin.

Flavored habit parameters. The parameters γ_2, γ_3, γ_4, and ψ_3 jointly govern the flavored habit stock and its effects on choice. Their identification rests on a key asymmetry: the nicotine addiction stock rises with any nicotine purchase, whereas the flavored habit stock rises only with flavored e-cigarette purchases. This means two households can have the same overall nicotine addiction level but very different flavored habit stocks depending on their recent purchasing histories, and that independent variation is the primary source of identification.

The lock-in penalty γ_2 for original (non-flavored) e-cigarettes and original bundles is identified from how the probability of choosing those alternatives varies with a household's recent flavored purchasing history, holding the nicotine addiction stock fixed. A household that bought flavored e-cigarettes every month for the past six months has a high flavored habit stock; if it is less likely to choose unflavored e-cigarettes than an otherwise identical household with a low flavored habit stock, this reveals $\gamma_2 < 0$. The cigarette lock-in penalty γ_3 is identified analogously: among households with similar nicotine addiction levels, those with a history of heavy flavored purchasing should be less likely to switch to cigarettes, and the magnitude of this effect pins down γ_3. The flavored withdrawal cost γ_4 is identified from the outside option: conditional on nicotine addiction (which is already captured by γ_1), households with a high flavored habit stock should be less likely to exit the market entirely. The separate variation in nicotine addiction and flavored habit stocks across households allows γ_1 and γ_4 to be separately identified: γ_1 is driven by any purchase history that builds nicotine addiction, while γ_4 is driven specifically by flavored e-cigarette purchase history.

The habit decay rate ψ_3 is identified from how quickly a household's elevated probability of returning to flavored products fades after it stops buying them. If the return probability collapses within a month or two of stopping, the habit stock must be depleting rapidly, implying a large ψ_3. If instead households that have been away from flavored products for several months still return at elevated rates relative to those absent for much longer, the stock is persisting, implying a small ψ_3. This variation in the speed of return across households with different gap lengths pins down the depreciation rate.

Type-specific fixed effects, mixing weights, and unobserved heterogeneity. A central identification challenge in any demand model is unobserved heterogeneity: consumers differ in ways the econometrician cannot see, and if those differences are correlated with the regressors of interest, parameter estimates will be biased. In the data, some households are lifelong cigarette buyers who never touch an e-cigarette, while others buy e-cigarettes almost exclusively. These persistent category loyalties survive even after accounting for prices, flavors, and addiction, pointing to deeper time-invariant differences in preferences that are not captured by any observed covariate. A single set of category fixed effects cannot fit both patterns simultaneously: any ξ_E large enough to explain the

e-cigarette loyalists would overpredict e-cigarette demand among the cigarette loyalists, and vice versa. The type-specific fixed effects ξ_{lC}, ξ_{lE}, and ξ_{lCE} for $l \in \{1, 2\}$ resolve this by allowing each latent type to have its own baseline category valuations, effectively acting as household fixed effects for time-invariant unobserved preference heterogeneity. The mixture likelihood sorts households into types based on their full observed purchasing histories, so that the estimated fixed effects absorb the portion of category choice driven by unobserved tastes rather than by prices, addiction, or flavor. The mixing weight parameters π_0 and π_{TYA} are identified from the relationship between these persistent patterns and household composition: if households with teens or young adults are disproportionately of the type that favors e-cigarettes, this correlation pins down π_{TYA}.

VI.2. *Independence of Irrelevant Alternatives*

The type I extreme value assumption on ε_{iljt} generates the logit form for conditional choice probabilities, which carries with it the independence of irrelevant alternatives (IIA) property: conditional on type and state, removing an alternative redistributes its demand to all remaining alternatives in exact proportion to their current shares. In a simple logit, this means a flavor ban would send displaced e-cigarette demand to cigarettes, unflavored e-cigarettes, and the outside option in fixed proportions regardless of who the consumer is or what their history looks like. For this paper's counterfactuals, where the central question is *how* demand reallocates following a ban, relying on this mechanical proportionality would be a serious problem. Three features of the model break it.

First, the two latent types have different baseline category preferences captured by $\xi_{lk(j)}$. Type 1, which has a strong baseline preference for cigarettes, responds to a flavor ban primarily by substituting toward cigarettes. Type 2, which has a stronger baseline preference for e-cigarettes, responds primarily by shifting to unflavored e-cigarettes. At the population level, the overall substitution pattern is a mix of these two very different responses, weighted by the share of each type, and so is not proportional in the way a homogeneous logit would predict.

Second, the nicotine addiction stocks generate heterogeneous substitution across households. A heavily addicted household faces a large abstention cost $\gamma_1 \cdot a_{it}$ if it exits the market, so its displaced demand flows toward other nicotine products rather than toward non-purchase. A lightly addicted household faces little abstention cost and may simply stop buying. Since addiction stocks vary across households and evolve with past choices, the model produces substitution patterns that depend on individual histories rather than population shares.

Third, the flavored habit stock makes substitution depend on a household's specific flavored purchasing history. Through γ_2, households with entrenched flavored habits face extra resistance to switching to unflavored e-cigarettes; through γ_3, they face separate

resistance to switching to cigarettes; and through γ_4, they face a higher cost of exiting the market entirely. These three penalties vary independently across households and allow the model to separately determine how much counterfactual demand flows toward each of the available alternatives.

VI.3. *Transition Probabilities*

The state vector $\boldsymbol{x}_{it} = \left(h_{it}, a^f_{it}, a^s_{it}, a^{\text{flav}}_{it}, p_{iCt}, p_{iEt}\right)$ contains four types of components: the binary teen-or-young-adult indicator h_{it}, which is a fixed household covariate; the nicotine addiction stocks a^f_{it} and a^s_{it} and the flavored habit stock a^{flav}_{it}, all of which evolve deterministically conditional on choices; and prices p_{iCt} and p_{iEt}, which follow a stochastic process. The nicotine decay rates ψ_2 and ψ_1 and the AR(1) price process parameters are estimated or calibrated in a first stage and held fixed during the second-stage maximization over $\boldsymbol{\theta}^u$. The first-stage transition parameters are $\boldsymbol{\theta}^f = (\psi_2, \psi_1, \phi_{0C}, \phi_{0E}, \phi_{1C}, \phi_{1E}, \Sigma)$. The flavor habit decay rate ψ_3 is estimated jointly with the structural utility parameters in the second stage.

Addiction. Given a household's choice in period t, both addiction stocks update deterministically:

$$a^s_{i,t+1} = (1-\psi_1)\cdot a^s_{it} + n_{iCt} + n_{iEt},$$
$$a^f_{i,t+1} = (1-\psi_2)\cdot a^f_{it} + n_{iCt} + n_{iEt},$$

where $\psi_1 = 0.10$ and $\psi_2 = 0.90$ are calibrated from the biomedical literature as discussed in Section V.4. and Section X.6.. The law of motion is deterministic conditional on choice, so no integration over addiction states is required when computing continuation values.[42]

Flavored habit stock. The flavored habit stock also updates deterministically given the household's choice:

$$a^{\text{flav}}_{i,t+1} = (1-\psi_3)\cdot a^{\text{flav}}_{it} + \mathbb{1}\left\{j_{it} \in \mathcal{E}_f\right\},$$

where $\psi_3 \in (0,1)$ is the flavor habit decay rate, estimated jointly with the other structural parameters in the second stage. Unlike the nicotine decay rates, which are pinned down by calibration before estimation, ψ_3 is a free parameter and its value shapes how persistent the flavor-specific inertia is. Since the law of motion is deterministic con-

[42] Although only the average $a_{it} = (a^f_{it} + a^s_{it})/2$ enters the flow utility, both stocks must be tracked separately as state variables. For example, take two households each with an average of 0.5, but the first household has $a_f = 0.8$ and $a_s = 0.2$, while the second has $a_f = 0.2$ and $a_s = 0.8$. Suppose also that they do not consume anything this period. This implies next period a_f for both households is 0.02 and 0.08, respectively, while next period a_s for both households is 0.72 and 0.18. Taking the average for each household gives next period addiction of 0.37 and 0.13. Consequently, both households have identical starting average and identical consumption this period, but very different addiction states next period, so simply tracking the aggregation does not suffice when computing continuation values.

ditional on choice, the flavored habit stock requires no additional integration in the Bellman equation.

Prices. I estimate category-level AR(1) processes on monthly median per-unit prices:

$$p_{C,t+1} = \phi_{0C} + \phi_{1C} \cdot p_{Ct} + \eta_{C,t+1},$$
$$p_{E,t+1} = \phi_{0E} + \phi_{1E} \cdot p_{Et} + \eta_{E,t+1},$$

with correlated shocks to these prices taking the form

$$\begin{bmatrix} \eta_{C,t+1} \\ \eta_{E,t+1} \end{bmatrix} \sim \mathbb{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma \right).$$

Unlike addiction, prices are stochastic from the consumer's perspective: a household knows current prices but not next month's price shocks. To integrate over future prices in the Bellman equation, I use the estimated coefficients $\left(\widehat{\phi}_{0C}, \widehat{\phi}_{1C}, \widehat{\phi}_{0E}, \widehat{\phi}_{1E}, \widehat{\Sigma}\right)$ to generate correlated price draws via quasi-Monte Carlo Halton sequences for every current price-grid point, and average over the resulting continuation values.

Teen or young-adult state. I classify each household-month into a binary demographic state $h_{it} \in \{0, 1\}$, where $h_{it} = 1$ indicates that at least one household member is aged 13–25 in period t and $h_{it} = 0$ otherwise. Since the teen-or-young-adult indicator is a predetermined household characteristic that does not respond to nicotine purchasing decisions, I treat h_{it} as a fixed covariate that shifts flow utility but does not enter the continuation value. This avoids the need to model transitions in household composition and eliminates a state dimension from the dynamic programming problem.

VI.4. *Estimation Framework*

Consumers $i \in \{1, 2, \ldots, N\}$ solve an infinite-horizon dynamic problem with exponential discounting ($\beta = 1$, $\delta = 0.99$).[43] At period t, consumer i observes the state $\boldsymbol{s}_{it} = (\boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it})$, where $\boldsymbol{x}_{it} = \left(h_{it}, a^f_{it}, a^s_{it}, a^{\text{flav}}_{it}, \boldsymbol{p}_{it}\right) \in \mathcal{X}$ is the observable state space and $\boldsymbol{\varepsilon}_{it} \in \mathbb{R}^{|\mathcal{J}|}$ is the vector of idiosyncratic i.i.d. type-one extreme value shocks for each alternative. Then, each consumer selects one alternative $j \in \mathcal{J} = \{1, \ldots, |\mathcal{J}|\}$ and the choice-specific value for this alternative j given type l is

$$\overline{v}_{iljt}\left(\boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \widehat{\boldsymbol{\theta}^f}\right) = \overline{u}_{iljt}\left(\boldsymbol{x}_{it}; \boldsymbol{\theta}^u\right) + \delta \int_{\mathcal{X}} V_l\left(\boldsymbol{x}_{i,t+1}; \boldsymbol{\theta}^u, \widehat{\boldsymbol{\theta}^f}\right) dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \widehat{\boldsymbol{\theta}^f}\right),$$

[43] The infinite-horizon formulation is standard in addiction models (e.g., Becker and Murphy (1988)) and avoids imposing an arbitrary terminal age. I estimate the model under standard exponential discounting and then explore how the counterfactual flavor ban results change as β varies from 1 in Section VIII.. This approach follows Gruber and Köszegi (2001), who show that the welfare implications of addiction policy depend critically on the degree of present bias, and allows me to trace out the sensitivity of the policy conclusions without taking a stand on a single value of β during estimation.

where $\overline{u}_{iljt}$ is the deterministic component of the flow utility for type l and V_l is the type-specific continuation value. Because each latent type has its own category fixed effects $\xi_{lk(j)}$, the flow utilities, value functions, and choice probabilities are all type-specific. The continuation value for type l satisfies the Bellman equation

$$V_l\left(\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\widehat{\boldsymbol{\theta}^f}\right)=\log\left(\sum_{j=1}^{|\mathcal{J}|}\exp\left[\overline{v}_{iljt}\left(\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\widehat{\boldsymbol{\theta}^f}\right)\right]\right),$$

which follows from the type one extreme value assumption on ε_{iljt}.[44] The conditional choice probability for alternative j given type l and observed state $\boldsymbol{x}_{it}$ is

$$\mathbb{P}_{iljt}\left(\boldsymbol{x}_{it}\right)=\frac{\exp\left[\overline{v}_{iljt}\left(\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\widehat{\boldsymbol{\theta}^f}\right)\right]}{\sum_{j'=1}^{|\mathcal{J}|}\exp\left[\overline{v}_{ilj't}\left(\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\widehat{\boldsymbol{\theta}^f}\right)\right]}.$$

The key computational implication is that each evaluation of the objective function requires solving the Bellman equation for both types. I compute this by value function iteration over the discretized state space, solving both types in parallel, and evaluate the continuation integrals by Monte Carlo integration. (see Section X.4.).

Mixture likelihood. Since the household's type l is unobserved, the full likelihood for household i integrates over both types. The finite mixture is the dynamic model analogue of a random coefficients specification: both allow preferences to vary across consumers, but the finite mixture places mass at discrete points rather than integrating over a continuous distribution.[45] Type membership is not imposed by me, but rather the likelihood assigns each household a posterior probability of belonging to each type based on its observed purchasing history. The prior mixing weight for type 2 is

$$\pi_{i2}=\frac{e^{\pi_0+\pi_{\text{TYA}}\cdot\overline{h}_i}}{1+e^{\pi_0+\pi_{\text{TYA}}\cdot\overline{h}_i}},\qquad\pi_{i1}=1-\pi_{i2},$$

where π_0 is a baseline intercept, π_{TYA} is a TYA share shifter, and $\overline{h}_i$ is the household's time-averaged share of months with a teen or young adult present.[46] Parameterizing the mixing weight as a function of TYA composition allows the type distribution to shift with household composition, enabling the model to rationalize the differences in

[44] The assumptions required for this representation are stated in Section X.3..

[45] In a static model, random coefficients are manageable via simulation over draws from the mixing distribution. In a dynamic model, however, each draw requires a separate VFI solve since the value function depends on the consumer-specific parameters, making a continuous distribution computationally prohibitive. The finite mixture requires only $L = 2$ VFI solves per likelihood evaluation regardless of sample size.

[46] The mixing weight uses the time-averaged $\overline{h}_i$ rather than the period-specific h_{it} because latent type is a permanent household characteristic, not a time-varying one. Using h_{it} would allow a household's type probability to change each period as TYA status changes, which is inconsistent with the interpretation of types as fixed unobserved heterogeneity.

category choice patterns across TYA and non-TYA households.

The full likelihood for household i is

$$\mathcal{L}_i\left(\boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) = \sum_{l=1}^{2} \pi_{il} \cdot \mathbb{P}\left(\boldsymbol{x}_{i1} \mid \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) \cdot \prod_{t=1}^{T_i} \mathbb{P}_{iljt}\left(\boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) \cdot \prod_{t=1}^{T_i-1} f_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid j_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right).$$

Inside the sum over types, the three components are: (i) the density of the initial state $\mathbb{P}(\boldsymbol{x}_{i1})$, (ii) the product of type-specific conditional choice probabilities across periods, and (iii) the product of state transition densities. Following the standard two-stage simplification, the initial-conditions density does not depend on $\boldsymbol{\theta}^u$ and drops out, and the transition densities are fully determined by the first-stage estimates $\widehat{\boldsymbol{\theta}^f}$ and are therefore constant with respect to $\boldsymbol{\theta}^u$. In my model this is natural: addiction and flavor habit evolves deterministically conditional on choices, and the price process is pre-estimated in stage one. What remains is a pseudo-likelihood based solely on the type-specific conditional choice probabilities:

$$\max_{\boldsymbol{\theta}^u} \mathscr{L}\left(\boldsymbol{\theta}^u, \widehat{\boldsymbol{\theta}^f}\right) = \max_{\boldsymbol{\theta}^u} \left\{\sum_{i=1}^{N} \log\left[\sum_{l=1}^{2} \pi_{il} \cdot \prod_{t=1}^{T_i} \mathbb{P}_{iljt}\left(\boldsymbol{x}_{it}\right)\right]\right\},$$

where $\mathscr{L}$ denotes the second-stage pseudo-log-likelihood. The inner product $\prod_{t=1}^{T_i} \mathbb{P}_{iljt}$ is the probability of observing household i's entire purchase history if it were type l. The inner sum integrates over the unobserved type by weighting each type's panel likelihood by the household-specific mixing probability π_{il}. The model never hard-assigns a household to a single type as a household whose purchasing pattern is much better explained by one type's parameters will naturally have most of its likelihood mass concentrated on that type, while a household equally well explained by both contributes more evenly.

Two-stage procedure. The numerical procedure is a two-stage nested fixed-point routine. In the first stage, I estimate or calibrate the transition parameters $\widehat{\boldsymbol{\theta}^f} = \left(\widehat{\psi}_f, \widehat{\psi}_s, \widehat{\boldsymbol{\phi}}_{0k}, \widehat{\boldsymbol{\phi}}_{1k}, \widehat{\Sigma}\right)$, which consist of the calibrated nicotine addiction decay rates, the AR(1) price process coefficients, and the price shock covariance matrix. In the second stage, I condition on $\widehat{\boldsymbol{\theta}^f}$ and maximize $\mathscr{L}\left(\boldsymbol{\theta}^u, \widehat{\boldsymbol{\theta}^f}\right)$ over the structural utility parameters $\boldsymbol{\theta}^u$. The parameter vector $\boldsymbol{\theta}^u$ includes the flavor habit decay rate ψ_3 in addition to the preference parameters, so the habit stock transitions, initial stock computations, and simulated habit trajectories are all recomputed at each candidate ψ_3 inside the second-stage objective. Each evaluation of the objective function at a candidate $\boldsymbol{\theta}^u_r$ requires: (i) solving the Bellman equation for both types by iterating until $\left\|V_l^{(n+1)} - V_l^{(n)}\right\|_\infty < \tau$ for each $l \in \{1, 2\}$; (ii) constructing the type-specific choice probabilities $\mathbb{P}_{iljt}$; and (iii) evaluating the mixture pseudo-log-likelihood $\mathscr{L}$. The optimizer then updates the candidate parameter vector and the cycle repeats until convergence. I implement this outer

optimization with a derivative-free quasi-Nelder-Mead method, which is well-suited to this setting because computing gradients of the value function is computationally impractical.[47]

After obtaining $\widehat{\boldsymbol{\theta}^u}$, I compute standard errors from the inverse numerical Hessian of the objective function (details in Section X.4.12.). Full mathematical details, assumptions, and computational implementation are provided in Section X.3. and Section X.4..

VII. Model Results

VII.1. *Estimates*

Consumption utility. Table 4 reports the estimated structural parameters. Cigarette and e-cigarette consumption both enter utility positively ($\widehat{\alpha}_C = 0.0239$, $\widehat{\alpha}_E = 0.0098$). Since α_C is in utils per pack and α_E is in utils per mL, the two coefficients are not directly comparable. Using the nicotine-absorption equivalence derived in Section X.7., 1 mL of e-liquid at a standard 50 mg/mL concentration delivers approximately the same absorbed nicotine as one pack of cigarettes. Thus, cigarettes generate roughly $0.0239/0.0098 \approx 2.4$ times the consumption utility of 1 mL of e-liquid, indicating that cigarette utility extends well beyond nicotine delivery alone. These magnitudes scale with quantity: a household buying the median 10 packs of cigarettes receives a consumption utility of $0.0239 \times 10 = 0.239$, while one purchasing the median 15.2 mL of e-cigarette liquid receives $0.0098 \times 15.2 = 0.149$. The negative bundle interaction ($\widehat{\alpha}_{CE} = -0.0014$) confirms that cigarettes and e-cigarettes are within-period substitutes, so households buying both products experience diminished marginal utility from each.

Flavor preferences. The flavor parameters reveal a large and demographically differentiated preference for flavored products. The baseline flavor premium is $\widehat{\lambda}_1 = 0.1529$, and it nearly doubles for households with a teen or young adult present: the combined TYA flavor effect is $\widehat{\lambda}_1 + \widehat{\lambda}_2 = 0.3521$, roughly twice the non-TYA value. This heterogeneity is central to the policy question, since the households most drawn to flavored products are precisely those flavor restrictions aim to protect. Among flavored products, households slightly prefer unauthorized over FDA-authorized brands $\left(\widehat{\lambda}_3 = -0.0627\right)$, a preference especially pronounced among TYA households $\left(\widehat{\lambda}_4 = -0.4412\right)$.

[47] For further details, see Section X.4.11..

Table 4. Structural Parameter Estimates.

Description	Parameter	Coefficient	Std. Error
Consumption Utility			
Cigarette consumption	$\widehat{\alpha}_C$	0.0239	(0.0007)
E-cigarette consumption	$\widehat{\alpha}_E$	0.0098	(0.0001)
Bundle interaction	$\widehat{\alpha}_{CE}$	−0.0014	(0.0000)
Flavor Preferences			
Flavored E-cig.	$\widehat{\lambda}_1$	0.1529	(0.0042)
Flavored E-cig. × TYA	$\widehat{\lambda}_2$	0.1992	(0.0058)
FDA Flavored E-Cig.	$\widehat{\lambda}_3$	−0.0627	(0.0111)
FDA Flavored E-Cig. × TYA	$\widehat{\lambda}_4$	−0.4412	(0.0139)
Addiction Dynamics			
Abstention cost	$\widehat{\gamma}_1$	−0.007390	(0.000031)
E-cig./bundle lock-in penalty	$\widehat{\gamma}_2$	−5.2151	(2.1948)
Cigarette lock-in penalty	$\widehat{\gamma}_3$	−1.0791	(0.4360)
Flavored withdrawal cost	$\widehat{\gamma}_4$	−0.3623	(0.1504)
Flavored habit decay rate	$\widehat{\psi}_3$	0.7401	(0.1013)
Price Sensitivity			
Cigarette price	$\widehat{\omega}_C$	−0.000917	(0.000107)
E-cigarette price	$\widehat{\omega}_E$	−0.001087	(0.000036)
Category Fixed Effects, Type 1			
Cigarettes	$\widehat{\xi}_{1C}$	−2.8193	(0.0049)
E-cigarettes	$\widehat{\xi}_{1E}$	−8.7898	(0.0802)
Bundles	$\widehat{\xi}_{1CE}$	−6.0555	(0.0465)
Category Fixed Effects, Type 2			
Cigarettes	$\widehat{\xi}_{2C}$	−4.1675	(0.0572)
E-cigarettes	$\widehat{\xi}_{2E}$	−3.3591	(0.0137)
Bundles	$\widehat{\xi}_{2CE}$	−2.8780	(0.0287)
Unobserved Heterogeneity			
$\mathbb{P}$(Type 1 | TYA Never Present)	$\widehat{\pi}$	0.9018	(0.0034)
$\mathbb{P}$(Type 1 | TYA Always Present)	$\widehat{\pi} + \widehat{\pi}_{\text{TYA}}$	0.8689	(0.0044)

Standard errors computed from the inverse numerical Hessian of the negative log-likelihood via central finite differences. The reported mixing probabilities are $\mathbb{P}(\textit{Type 1} \mid \bar{h} = 0) = 1 - \sigma(\hat{\pi}_0)$ *and* $\mathbb{P}(\textit{Type 1} \mid \bar{h} = 1) = 1 - \sigma(\hat{\pi}_0 + \hat{\pi}_{TYA})$*, where* $\sigma(x) = e^x/(1+e^x)$ *is the logistic function; standard errors are obtained by the delta method applied to this nonlinear transform. Standard errors for* γ_2*,* γ_3*, and* γ_4 *are likewise obtained by the delta method and account for the covariance between each* γ *and* $\hat{\psi}_3$*. Details on both delta method calculations are in Section X.4.12.. Units:* α_C *in utils/pack,* α_E *in utils/mL,* α_{CE} *in utils/(pack* × *mL),* ω_C *in utils/(\$* × *pack),* ω_E *in utils/(\$* × *mL),* γ_1 *in utils/mg of nicotine consumed per month,* γ_2*,* γ_3*, and* γ_4 *in utils per unit of raw flavored habit stock.*

Addiction dynamics. The abstention cost ($\widehat{\gamma}_1 = -0.007390$) is negative, meaning that a larger nicotine addiction stock raises the disutility of not purchasing. At steady state, the addiction stock equals the household's monthly nicotine consumption in milligrams. A smoker purchasing the median quantity of 10 packs per month absorbs roughly 250 mg of nicotine and therefore faces an abstention disutility of $-0.007390 \times 250 = -1.85$ utils from not purchasing.[48] Higher consumption builds the stock, the stock raises the cost of abstaining, and continued purchasing follows, producing the self-reinforcing dynamic of rational addiction.

Flavored habit stock. The flavored habit stock parameters capture a second source of state dependence, operating through flavor-specific purchasing history independently of nicotine addiction. All three disutility parameters are negative, reflecting a lock-in mechanism: a high flavored habit stock depresses utility from non-flavored alternatives, keeping habitual flavored purchasers locked in to flavored e-cigarettes. The lock-in is strongest toward unflavored e-cigarettes ($\widehat{\gamma}_2 = -5.2151$), weaker toward combustible cigarettes ($\widehat{\gamma}_3 = -1.0791$), and weakest toward exiting entirely ($\widehat{\gamma}_4 = -0.3623$). Habitual flavored users therefore face the greatest resistance to switching within the e-cigarette category and a smaller, but still meaningful, barrier to exiting the market. The estimated decay rate $\widehat{\psi}_3 = 0.7401$ indicates that the flavored habit stock erodes rapidly once purchasing stops: a single month without a flavored purchase reduces the stock to $(1 - 0.7401) = 26\%$ of its prior level.[49]

Price sensitivity and type heterogeneity. Price sensitivity is negative for both categories ($\widehat{\omega}_C = -0.000917$, $\widehat{\omega}_E = -0.001087$). These enter flow utility as expenditure-style terms $\omega_C \cdot p_{iCt} \cdot q_{iCt}$ and $\omega_E \cdot p_{iEt} \cdot q_{iEt}$, so the disutility from a price increase scales with quantity purchased. At the median per-pack price of \$6.77 and the median purchase quantity of 10 packs, a typical cigarette purchase carries a price disutility of $-0.000917 \times 6.77 \times 10 = -0.0621$ utils. The category fixed effects sharply differentiate the two latent types. Type 1, comprising roughly 90% of non-TYA households and 87% of TYA households, has a relatively high cigarette fixed effect $\left(\widehat{\xi}_{1C} = -2.8193\right)$ and a very low e-cigarette fixed effect $\left(\widehat{\xi}_{1E} = -8.7898\right)$, indicating a strong baseline preference for combustible cigarettes. Type 2 shows the reverse: a lower cigarette fixed effect $\left(\widehat{\xi}_{2C} = -4.1675\right)$ paired with a substantially higher e-cigarette fixed effect $\left(\widehat{\xi}_{2E} = -3.3591\right)$, consistent with active vapers for whom flavor restrictions bind

[48] Using the nicotine-absorption equivalence in Section X.7., one pack of cigarettes delivers $20 \times 1.25 = 25$ mg of absorbed nicotine, where 1.25 mg is the per-cigarette absorption standard. A household purchasing 10 packs per month therefore absorbs $10 \times 25 = 250$ mg per month.

[49] The standard error on $\widehat{\psi}_3$ is relatively large (0.3013) and the estimate is significant at the 5% level but not the 1% level. ψ_3 is identified from the speed at which households return to flavored purchasing after a gap: a high ψ_3 means the habit stock decays quickly during the gap, so the model predicts a slow and gradual return; a low ψ_3 means the stock persists through the gap, so return rates stay high. Identifying this from the data requires observing households who stop flavored purchasing for one or more months and then resume. Such spells are relatively rare in the sample producing the larger standard error.

most directly. Having a teen or young adult in the household raises the probability of being Type 2 by about 3.3 percentage points, from 9.8% to 13.1%, consistent with the descriptive evidence linking younger household members to higher e-cigarette purchase rates.

VII.2. *Model Fit*

To assess model fit, I simulate one-hundred forward choice sequences from the estimated model, letting each household's addiction stocks and flavor habit stock evolve from simulated rather than observed choices. At each month, the model draws a purchase from the logit probabilities implied by the simulated addiction states, flavor habit stock, and the observed prices and TYA status. This is a stricter test than one-step-ahead prediction: if the addiction dynamics are misspecified, simulated trajectories drift from reality.

Figure 10 plots streak persistence curves for cigarettes and flavored e-cigarettes by TYA status, with 95% simulation confidence bands. A streak persistence curve asks: conditional on having purchased a product for k consecutive months, how likely is the household to purchase it again next month? A rising curve is the hallmark of habit formation (the longer a purchasing streak, the harder it becomes to break) and directly tests whether the model's addiction and flavored habit stock dynamics generate the right degree of lock-in. The simulated curves track the actual curves closely for both products and both TYA subgroups, with simulated rates falling within or near the confidence bands across most streak lengths. Confidence bands widen at longer streaks for flavored e-cigarettes, reflecting the much smaller sample of long flavored purchasing streaks in the data.

Figure 10. Streak Persistence Curves: Actual vs. Model-Predicted Continuation Rates.

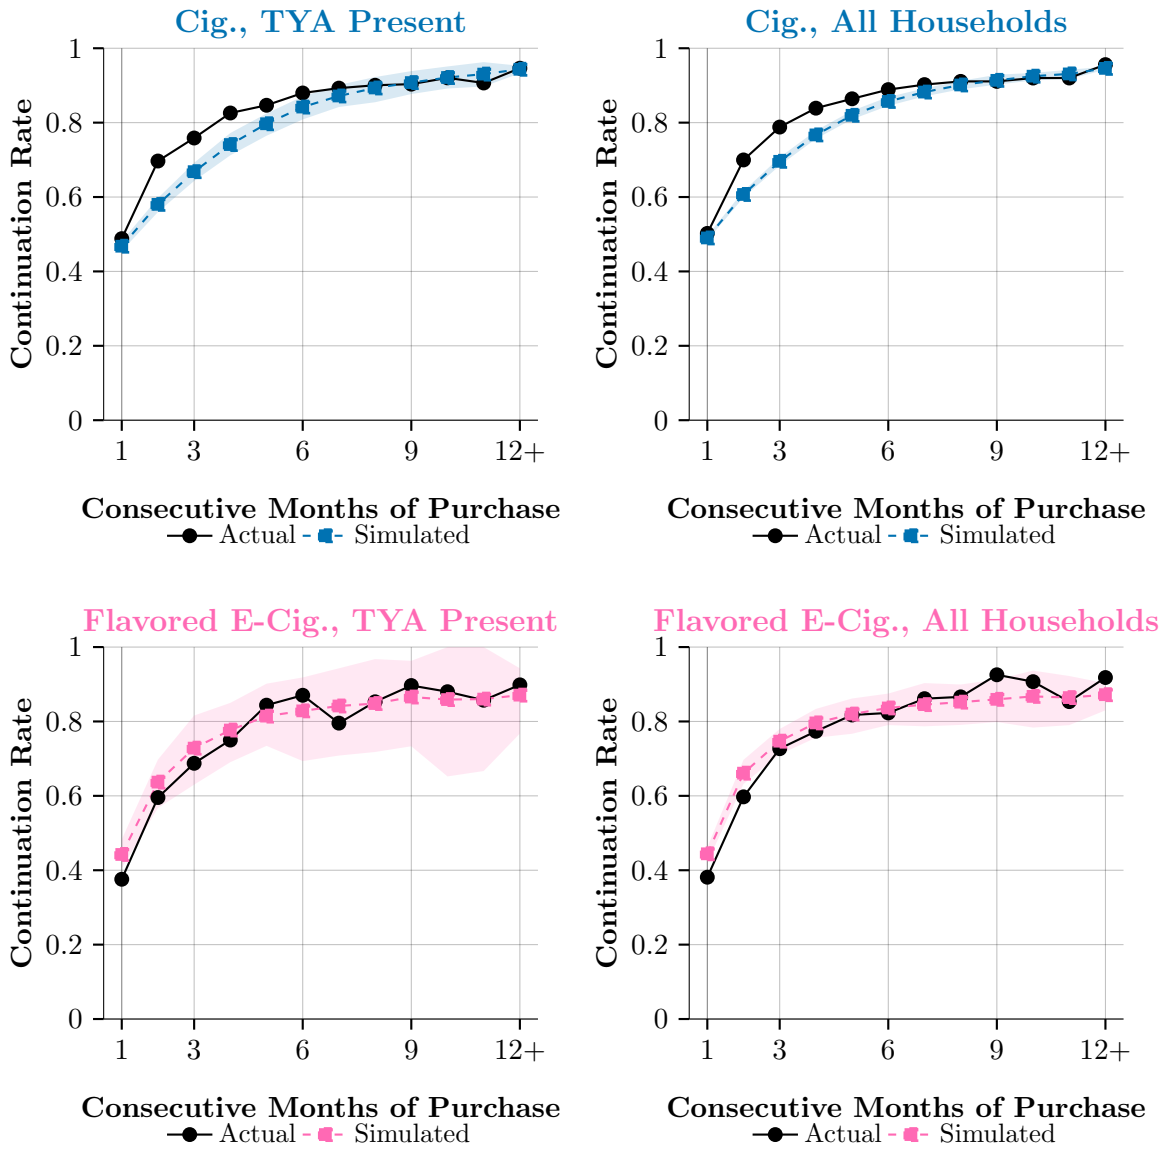

This figure compares actual and simulated streak continuation rates by consecutive months of purchase, separately for households with a teen or young adult present (left column) and all households (right column). The top row shows cigarettes and the bottom row shows flavored e-cigarettes. A streak of length k means the household purchased the product in each of the last k months. The continuation rate is the probability of purchasing again the following month. Actual rates are computed from observed choices. Simulated rates are the mean across $S = 100$ forward simulation draws in which the model generates choice sequences given the model's addiction dynamics; the shaded region shows the 95% confidence band. Streak lengths of 12 or more months are pooled.

As a second validation exercise, Figure 11 plots post-purchase paths following an initial flavored e-cigarette purchase: the share of households still purchasing flavored e-cigarettes and the share who have switched to cigarettes at each of the 12 subsequent months, separately by TYA status. This tests whether the model correctly captures

both the persistence of flavored purchasing and the direction of substitution after a flavored purchase event.

Figure 11. Post-Purchase Paths After a Flavored E-Cigarette Purchase.

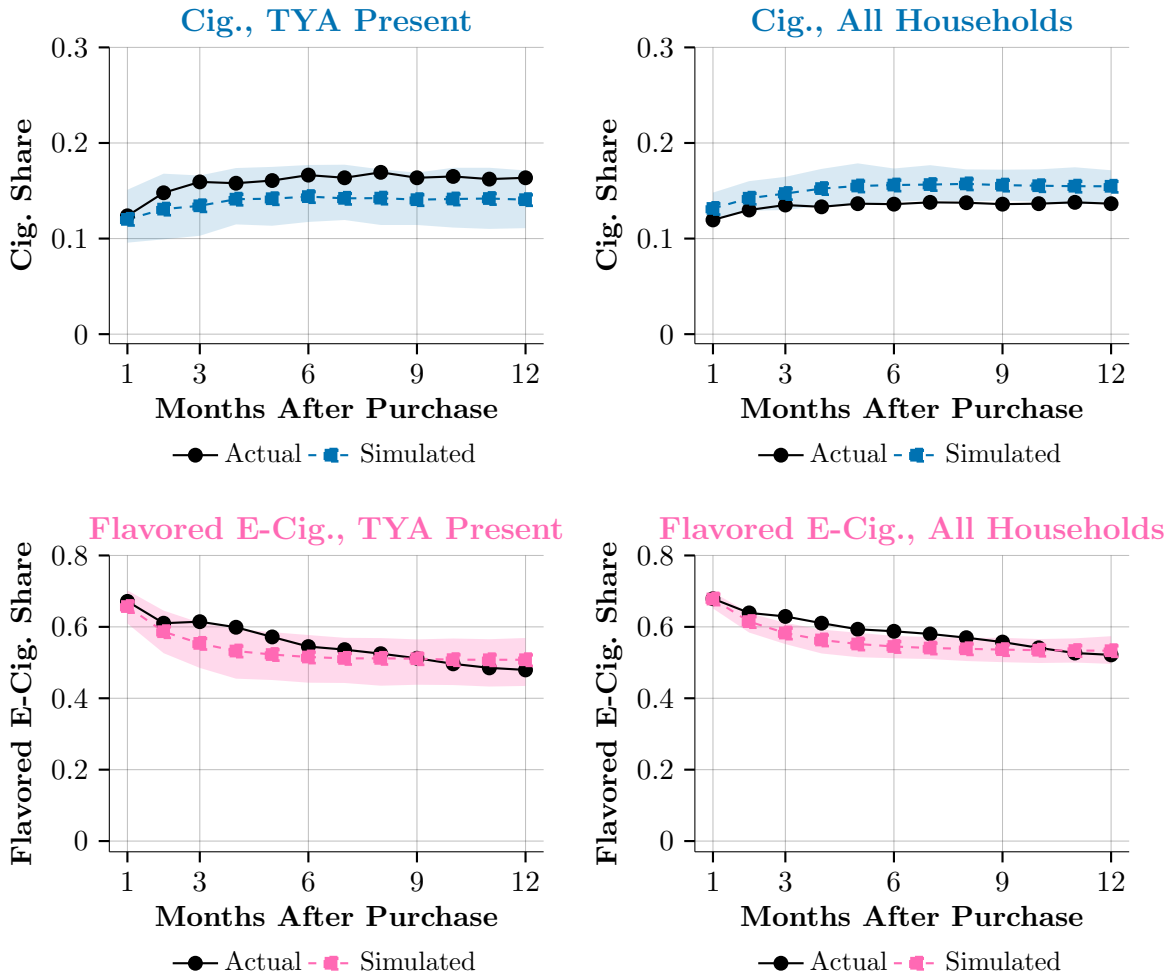

This figure plots category shares at each month following an initial flavored e-cigarette purchase, separately for TYA-present households (left column) and all households (right column). The top row shows cigarette shares and the bottom row shows flavored e-cigarette shares. Black circles show actual shares computed from observed choices. Colored dashed lines with squares show simulated means across $S = 100$ forward simulation draws in which the addiction and flavored habit stocks evolve endogenously based on simulated choices. Shaded bands show 95% simulation confidence intervals. Based on 612 TYA and 2,009 total flavored e-cigarette purchase events.

The simulated post-purchase paths track the actual paths closely for both products and both TYA subgroups. The model correctly replicates the gradual decline in flavored repurchasing over the 12-month horizon. Two features of the flavored post-purchase paths are particularly relevant for flavor policy. First, flavored purchasing declines over time following a flavored purchase event, as the flavored habit stock erodes without continued reinforcement and some households exit the nicotine market entirely. Second, cigarette purchasing rises modestly as the flavored share declines, capturing

the partial substitution of some households toward combustible cigarettes when their flavored habit weakens. Both patterns are important for evaluating flavor bans: they imply that removing flavored products not only reduces flavored purchasing directly but also triggers a dynamic reallocation, with most households exiting the nicotine market over time and a smaller share drifting toward cigarettes. Figure A1 in the Appendix performs the analogous exercise starting from a cigarette purchase and confirms that cigarette persistence is high in both the data and the model, while flows from cigarettes toward e-cigarettes are near zero throughout.[50]

VIII. Counterfactual Results

Three counterfactual policies are evaluated. The first removes all flavored e-cigarette products from the choice set, representing a perfectly enforced comprehensive ban. The second restricts the removal to FDA-authorized flavored products, capturing the realistic scenario in which unauthorized brands remain available. The third applies a per-milliliter excise tax on flavored e-cigarette liquid at three rates ($0.10, $0.25, and $0.50/mL). In each simulation, households begin from their end-of-sample addiction stocks rather than a fresh start, so the dynamics reflect the actual distribution of addiction at the time of the policy. Results are reported separately for TYA and non-TYA households and incorporate present bias ($\beta = 0.7$) to assess robustness to the degree of time inconsistency.[51]

VIII.1. *Banning All Flavored E-Cigarettes*

The first counterfactual simulates a comprehensive ban on all flavored e-cigarette products, removing both FDA-authorized and unauthorized flavored alternatives from the choice set. This scenario assumes perfect enforcement: every flavored e-cigarette and flavored bundle alternative receives flow utility $-\infty$, so the choice probability of any banned product is exactly zero. Consumers must substitute toward unflavored e-cigarettes, combustible cigarettes, bundles of cigarettes with unflavored e-cigarettes, or elect for the outside option.

Figure 12 plots the percentage decline in mean addiction under the comprehensive flavor ban relative to the status quo over 36 months, separately for TYA-present and

[50] Figure A1 shows cigarette persistence is consistently underpredicted by roughly ten percentage points, which likely occurs given there is not a cigarette habit stock analogous to the flavor habit stock already present in the model. However, the model does capture dynamic substitution patterns from cigarettes to e-cigarettes quite well, despite the small fraction of individuals making this switch. This is important for establishing how the model can evaluate how e-cigarette flavor policy impacts the extensive and intensive behavioral margins.

[51] The value $\beta = 0.7$ is consistent with, and in fact more conservative than, the calibration in the leading paper on time-inconsistent tobacco demand. Gruber and Köszegi (2001) calibrate $\beta = 0.6$ in their analysis of time-inconsistent cigarette demand and optimal sin taxes, and use a similar value in their follow-up work on the welfare effects of tobacco taxation Gruber and Köszegi (2004). Setting $\beta = 0.7$ therefore represents a milder degree of present bias than the benchmark in this literature. Results under both $\beta = 1$ and $\beta = 0.7$ are reported throughout.

TYA-absent households under both exponential ($\beta = 1$) and quasi-hyperbolic ($\beta < 1$) discounting. The comprehensive ban reduces mean addiction by 8.1% at 12 months and 12.2% at 36 months among time-consistent households ($\beta = 1$), with the reduction growing gradually as consumers who exit the nicotine market see their addiction stocks decay toward zero. TYA households experience a larger reduction: 10.8% at 12 months and 15.2% at 36 months, compared to 7.5% and 11.5% for households without youth members. The larger TYA effect reflects the higher estimated flavor premium for youth-exposed households ($\widehat{\lambda}_1 + \widehat{\lambda}_2 = 0.3521$), which means removing flavored products depresses the utility of remaining alternatives more sharply for this group and accelerates their exit from the market. The addiction reduction is robust to the degree of present bias: at $\beta = 0.7$ the 36-month reduction is 12.1%, nearly identical to 12.2% at $\beta = 1$.

Figure 12. Percentage Change in Mean Addiction by TYA Status and β.

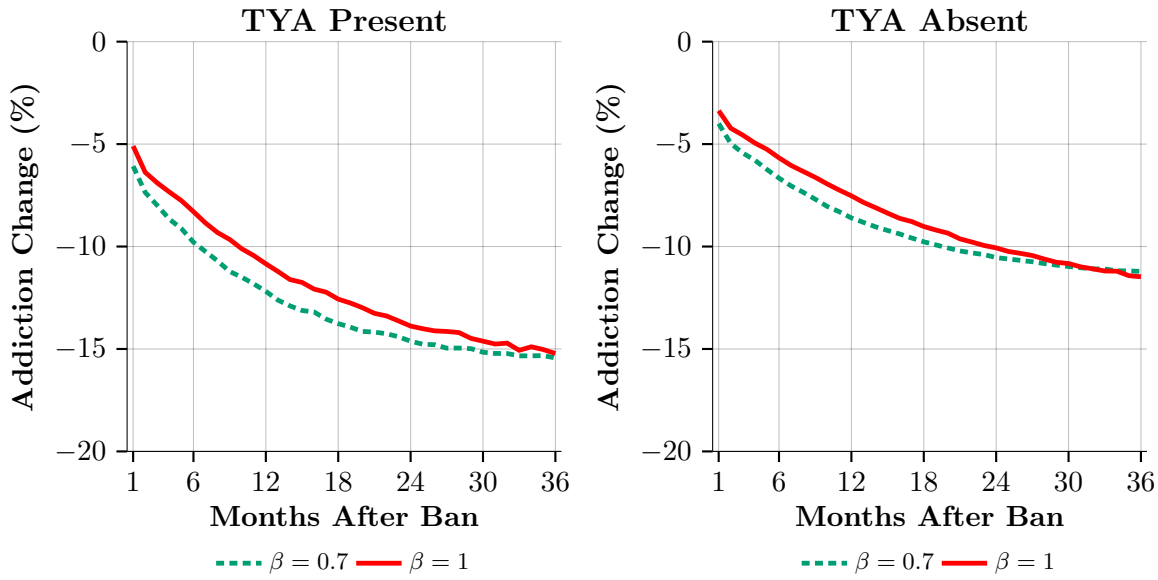

The key unintended-consequence concern, that banning flavored e-cigarettes would drive consumers toward combustible cigarettes, is short-lived. Cigarette market shares rise by roughly 0.7 percentage points at period 1, reflecting initial substitution, but this reverses by month 6 to 7: without flavored options to sustain the habit stock, consumers exit the nicotine market rather than settling into cigarette consumption, and by period 36 cigarette shares are below the status quo under both values of β. Original (unflavored) e-cigarettes absorb a modest share of displaced demand in the short run, with their market share rising by roughly 0.5 percentage points at $\beta = 1$ and 0.9 percentage points at $\beta = 0.7$ in the first few months after the ban. This substitution is also temporary: as the flavored habit stock erodes without reinforcement, even unflavored e-cigarettes

become less attractive, and the original e-cig share drifts back toward its status quo level by period 36. The larger short-run spike at $\beta = 0.7$ reflects the higher flavored consumption in the status quo under present bias; more consumers are displaced from a larger base, and some temporarily land on original e-cigs before ultimately exiting. Figure 13 plots both the cigarette and original e-cigarette share differences over the full 36-month horizon. The dominant margin of adjustment is exit.

Figure 13. Percentage Point Change in Cigarette and Original E-Cigarette Shares by TYA Status and β.

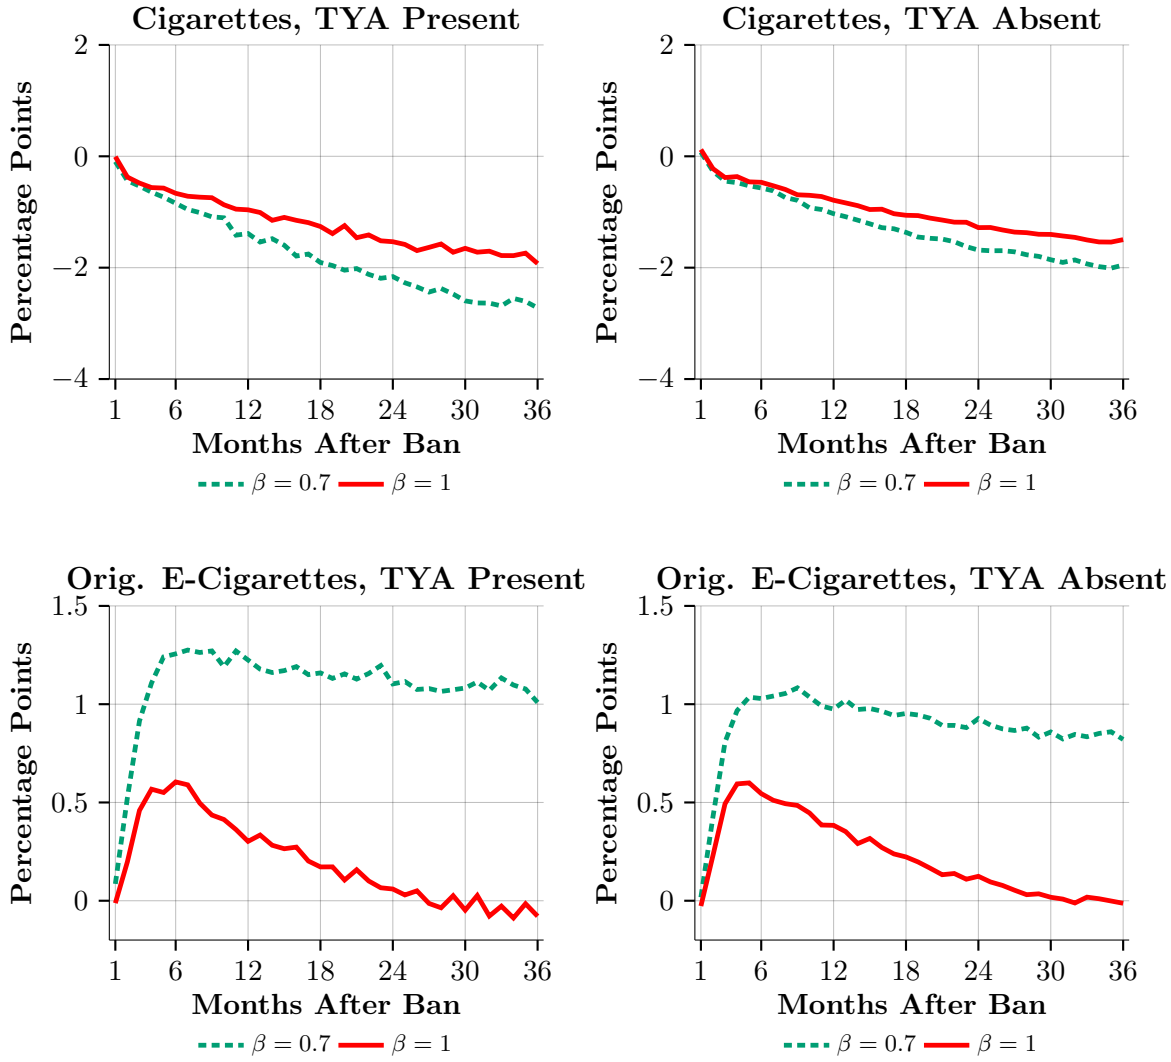

Under the ban the outside-option share exceeds the status quo by 1.7 percentage points in period 1 and by 3.9 percentage points at period 36 under $\beta = 1$: in the status quo, addiction accumulates over time and pulls consumers into the nicotine market, causing the outside-option share to decline steadily. The ban disrupts this process:

without flavored products to reinforce the habit stock, addiction builds more slowly and fewer consumers become locked in. Because the status quo continues drawing flavored e-cigarette consumers into the market each period while the ban does not, the outside-option gap widens over time. Figure 14 plots the full trajectories by TYA status and β; TYA households exit more rapidly, with a 4.7 percentage point gap at period 36 versus 3.6 percentage points for non-TYA households. Notably, the exit gap is larger under $\beta = 0.7$ than under $\beta = 1$ throughout the simulation. The reason is that present-biased consumers over-consume flavored products in the status quo, discounting future addiction costs too heavily, so the status quo outside-option share is lower at $\beta = 0.7$ than at $\beta = 1$ (57.3% versus 63.4% for TYA households at period 1). When the ban removes the flavored option that was sustaining this overconsumption, these households exit at higher rates precisely because the ban corrects the self-control failure that was keeping them in the market.

Figure 14. Percentage Point Change in Outside-Option Share by TYA Status and β.

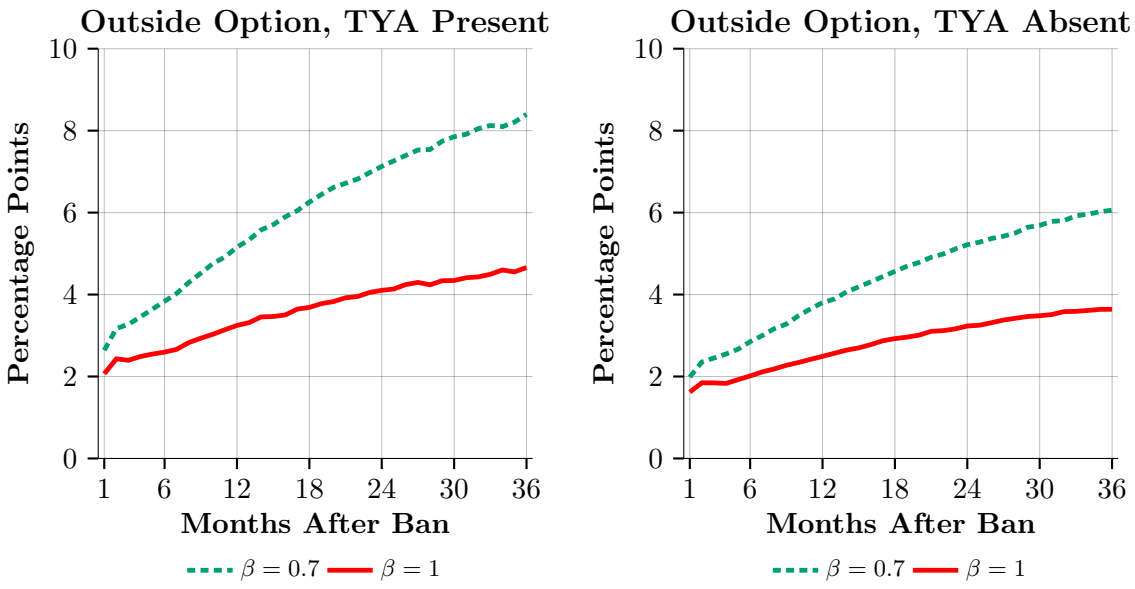

VIII.2. *Banning FDA-Authorized Flavored E-Cigarettes*

The preceding counterfactual assumes full enforcement across all flavored e-cigarette products, including those sold through unauthorized channels. In practice, however, an FDA-implemented flavor ban would only directly affect products that have received marketing authorization. As of 2026, only four e-cigarette brands (JUUL, Vuse, NJOY, and Logic) hold marketing orders from the FDA. The remaining products, which 2024 scanner estimates suggest account for roughly 86% of retail e-cigarette sales, are sold without authorization and would be largely unaffected by additional regulatory action

on flavors, as they already operate outside the legal market.

Comparing the comprehensive ban with the authorized-only ban provides a direct measure of the enforcement gap in U.S. e-cigarette regulation. The comprehensive ban is an upper bound on what a flavor ban can achieve under perfect enforcement; the authorized-only ban is a realistic lower bound, capturing what happens when only firms with regulatory exposure comply. The gap between the two quantifies how much the unauthorized market attenuates the intended effects of flavor regulation, and directly informs the value of investing in enforcement capacity.

As documented in Figure 5, the FDA-authorized share of flavored e-cigarette purchase months has declined substantially over the sample period and is consistently lower in households with a teen or young adult. Specifically, among flavored e-cigarette purchase months with households containing a teen or young adult, by 2023 roughly only 35% involve an FDA-authorized brand, compared to roughly 90% for unflavored purchases. This means that a realistic flavor ban targeting only authorized brands would directly affect less than half of flavored e-cigarette consumption, and would disproportionately miss the households with youth exposure to flavored nicotine products.

Figure 15 quantifies the enforcement gap along two margins. The top row plots the gap in addiction reduction between the comprehensive and FDA-only bans: negative values indicate that extending enforcement to unauthorized products would yield additional addiction reductions, with the gap reaching roughly 2.8 percentage points for TYA households by period 36. The bottom row shows the unauthorized flavored e-cigarette share that persists under the FDA-only ban (the flavored purchasing the comprehensive ban would eliminate but the FDA-only ban cannot), directly linking the residual substitution to unauthorized products with the addiction shortfall in the top row.

The enforcement gap is a central finding of this paper and speaks directly to the efficacy of the FDA's existing e-cigarette regulation. Because roughly 65% of flavored e-cigarette purchases by TYA households involve unauthorized brands by 2023 in my data, an FDA-only flavor ban is incapable of reaching the majority of flavored consumption in the households most at risk. The gap in addiction outcomes is not a second-order: the 2.8 percentage point TYA addiction shortfall at 36 months represents approximately 18% of the comprehensive ban's total addiction reduction for TYA households[52], and this share grows with present bias precisely because present-biased consumers are the most likely to substitute to unauthorized alternatives when authorized options are removed. The implication is stark: the population whose flavored e-cigarette habits the FDA most wants to curtail is the population least constrained by the instrument the FDA has authority to implement.

[52] Under $\beta = 1$, the comprehensive ban reduces mean TYA addiction by 15.2% at 36 months and the FDA-only ban reduces it by 12.4%, leaving a gap of 2.8 percentage points; $2.8/15.2 \approx 18\%$.

Figure 15. Enforcement Gap by TYA Status and β: Addiction Shortfall Relative to Comprehensive Ban (Top) and Residual Unauthorized Flavored Share Under FDA-Only Ban (Bottom).

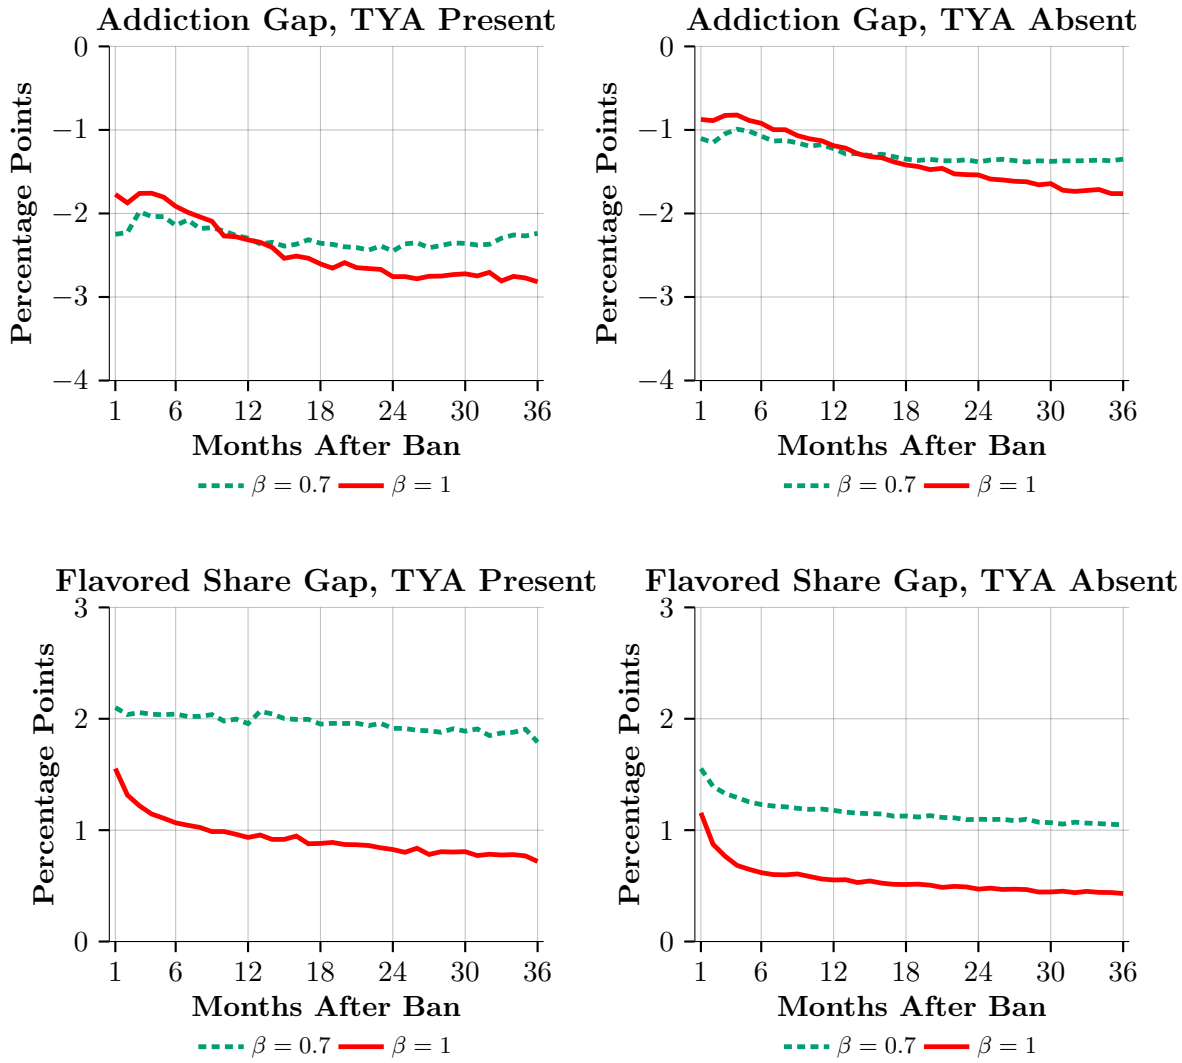

Figure 16 plots the percentage change in mean addiction under the FDA-only ban relative to the status quo, paralleling Figure 12 for the comprehensive ban. The FDA-only ban reduces mean addiction by 6.7% at 12 months and 10.2% at 36 months under $\beta = 1$, compared to 8.1% and 12.2% under the comprehensive ban. For TYA households the 36-month reduction is 12.4% (vs. 15.2% under the comprehensive ban), and for non-TYA households it is 9.7% (vs. 11.5%). The FDA-only ban thus achieves 84% of the total addiction reduction and 82% of the TYA-specific reduction, with the shortfall attributable to the unauthorized consumption that persists and keeps addiction stocks above the comprehensive ban's trajectory.

Figure 16. Percentage Change in Mean Addiction Under FDA-Only Ban Relative to Status Quo, by TYA Status and β.

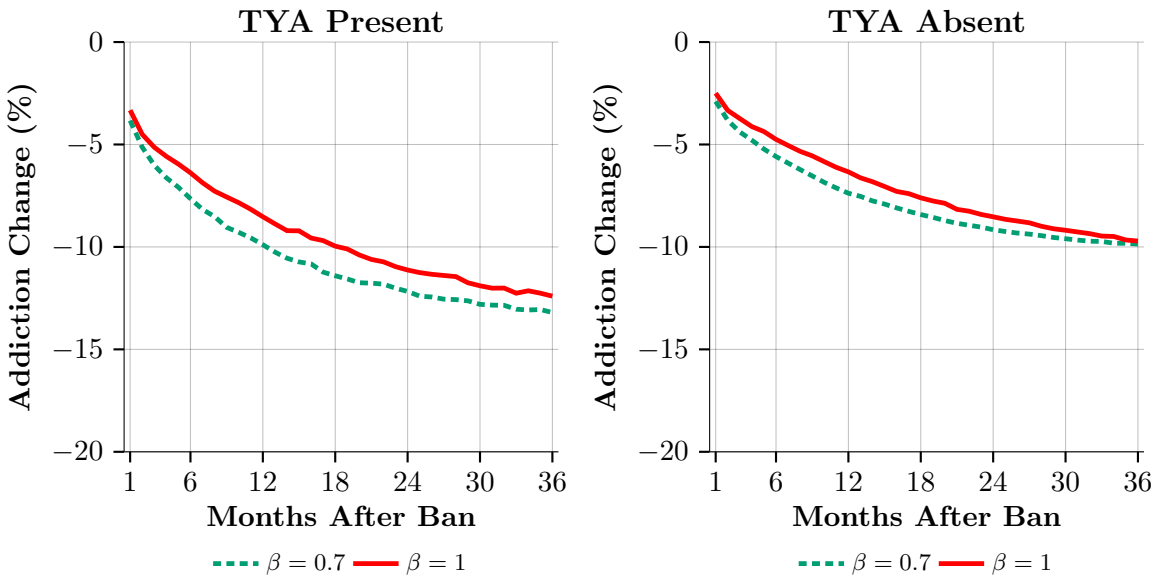

Figure 17 plots the percentage-point changes in unauthorized flavored e-cigarette and original e-cigarette market shares under the FDA-only ban relative to the status quo. Over the full 36-month horizon, unauthorized share falls well below the status quo (by as much as 3.6 percentage points for TYA households at $\beta = 0.7$) through an indirect channel. Removing authorized flavored products slows the rate at which the flavored habit stock is reinforced across the population; as the average habit stock decays, demand for unauthorized flavored alternatives erodes alongside it. The erosion is faster under $\beta = 1$, as time-consistent consumers are more likely to exit when their preferred authorized option disappears, and slower under $\beta = 0.7$, as present-biased consumers are more likely to persist with unauthorized substitutes. Original e-cigarettes absorb a small share of displaced demand, with the pattern differing by present-bias type: under $\beta = 1$, the original e-cigarette share rises modestly in the short run and fades back toward the status quo as habit stocks decay, while under $\beta = 0.7$ it rises persistently across the full 36 months. The mechanism is the internality: present-biased consumers previously over-consumed flavored products by discounting future addiction costs, entering the counterfactual with larger habit and addiction stocks. Those larger stocks raise the withdrawal penalty of abstention, so when authorized flavored products are removed, exiting to the outside option is more costly and substitution to unflavored alternatives is more attractive.

Figure 17. Percentage-Point Change in Category Shares Under FDA-Only Ban Relative to Status Quo, by TYA Status and β.

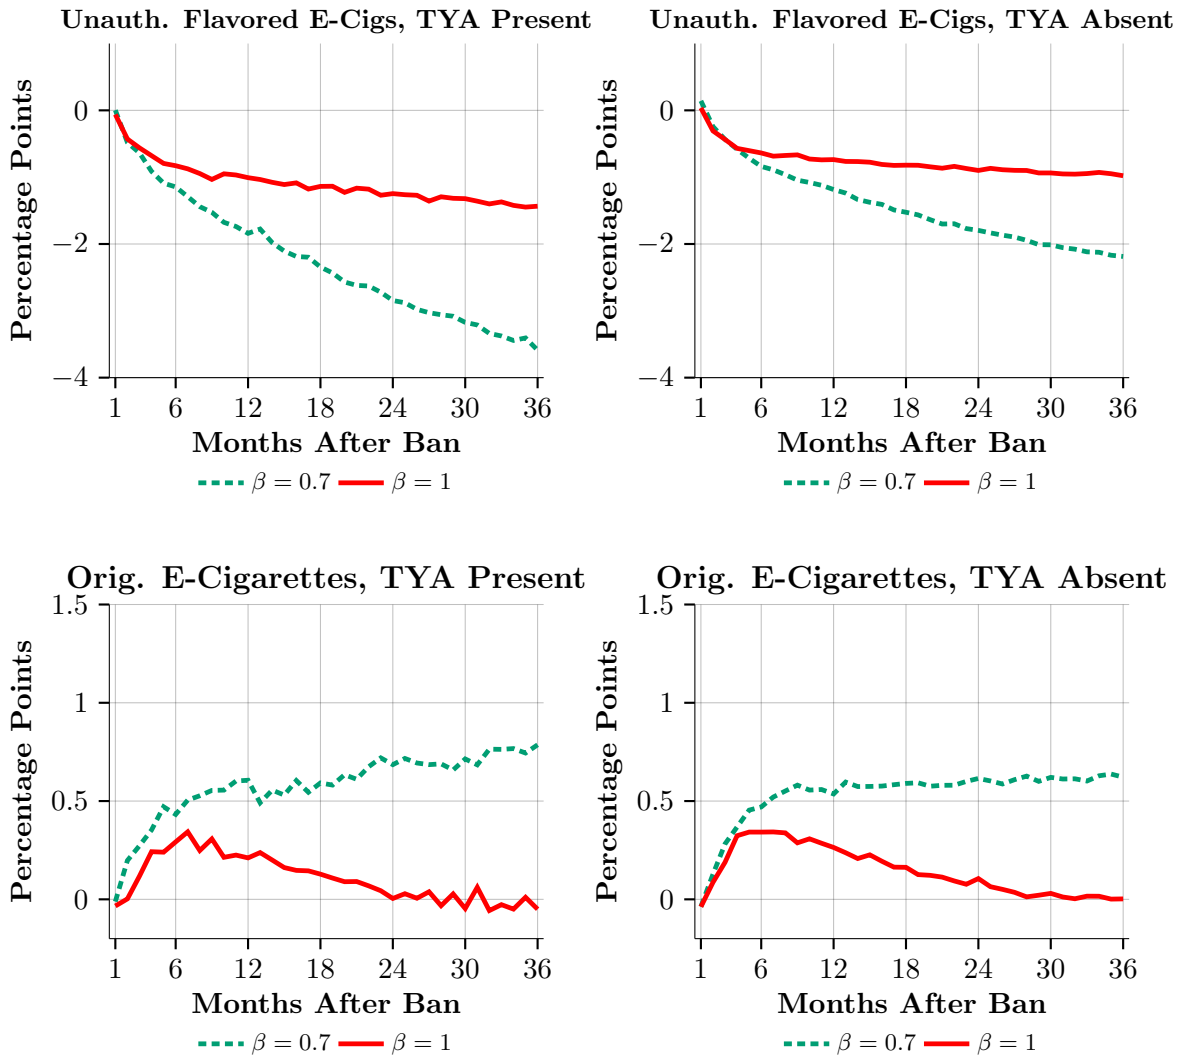

The FDA-only ban achieves meaningful reductions in addiction, but its impact falls short of the comprehensive ban precisely where the policy motivation is strongest. The 16–18% enforcement gap is driven by within-category diversion: consumers displaced from authorized brands substitute to unauthorized flavored alternatives with similar flavor profiles, reshuffling demand from regulated to unregulated sellers without proportionally constraining the overall availability of flavored nicotine.

VIII.3. *Flavored E-Cigarette Tax*

The preceding counterfactuals evaluate a quantity instrument. An alternative is a price instrument: a per-milliliter excise tax that raises the cost of flavored consumption without eliminating access. A ban forces substitution; a tax operates on the intensive

and extensive margins simultaneously, allowing consumers to reduce flavored consumption gradually while retaining the option to pay the tax and continue. The welfare implications also differ: a ban generates no revenue, while a tax generates revenue that could finance cessation programs or offset public health costs.[53]

Figure 18 plots the fraction of the comprehensive ban's addiction reduction achieved by the \$0.50/mL tax at each horizon, for both values of β. A value of 100% means the tax replicates the ban; values below 100% reflect consumers who continue purchasing flavored products at the taxed price rather than exiting. The tax converges toward 92–96% of the ban's addiction reduction by period 36, with the ratio rising over time as flavor habits decay. At $\beta = 1$ the ratio reaches 92% for TYA households (14.0% vs. 15.2%) and at $\beta = 0.7$ it reaches 96% (14.9% vs. 15.4%). The residual gap reflects consumers who continue purchasing at the taxed price rather than exiting, but it is small because the flavored habit stock generates strong lock-in: reducing flavored consumption erodes the stock quickly, lowers the future utility from flavored alternatives, and accelerates exit even under a price instrument.

Figure 18. Fraction of Comprehensive Ban's Addiction Reduction Achieved by \$0.50/mL Tax by TYA Status and β.

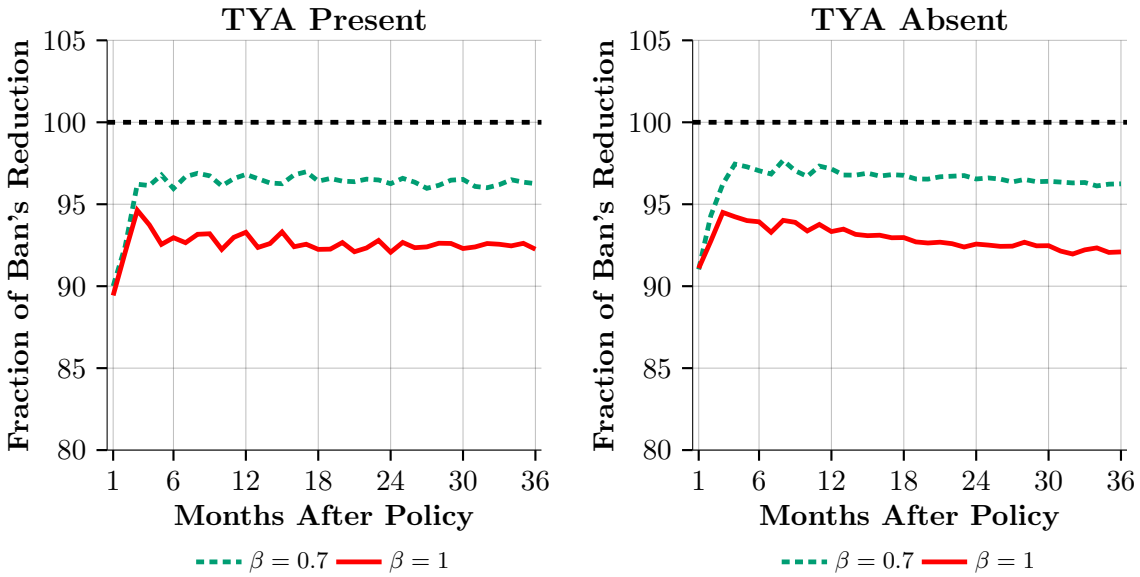

A practical advantage of the tax over the ban is that it generates revenue. Table 5 reports per-purchase revenue, an upper bound on annual U.S. revenue, and a demand-adjusted estimate that accounts for tax-induced reductions in purchase occasions. The upper bound scales the per-purchase tax by the number of annual purchase occasions

[53]The counterfactuals in this section assume 100% pass-through. This is a realistic representation given existing estimates suggest roughly 90% to 100% of the economic incidence of cigarette excise taxes are beared by consumers.

among current flavored e-cigarette users nationally. Drawing on two national surveys, I estimate approximately 5.5 million current flavored e-cigarette users in the U.S., a figure that is likely a lower bound given documented growth in the e-cigarette market since 2018.[54] Assuming each user purchases once per month yields 66 million annual occasions. The demand-adjusted estimate then scales this figure by the ratio of post-tax to pre-tax flavored purchase shares, averaged over the 36-month horizon at $\beta = 1$, capturing the progressive demand reduction induced by the tax.[55]

Table 5. Flavor Tax Revenue by Tax Rate for $\beta = 1$

Tax Rate	**Per Purchase**	**Upper Bound**	**Demand-Adjusted**
All-Flavored Tax			
\$0.10/mL	\$1.59	\$105M	\$46M
\$0.25/mL	\$3.98	\$263M	\$76M
\$0.50/mL	\$7.95	\$525M	\$98M
FDA-Only Flavor Tax			
\$0.10/mL	\$1.68	\$18M	\$9M
\$0.25/mL	\$4.20	\$44M	\$15M
\$0.50/mL	\$8.40	\$89M	\$20M

Per-purchase revenue is $\tau \times q_{med}$, where q_{med} is the median quantity across flavored quantity bins: 15.9 mL for the all-flavored tax (median of all 18 flavored e-cigarette and bundle bins) and 16.8 mL for the FDA-only tax (median of the 9 FDA-authorized bins). The upper bound assumes 5.5 million current flavored e-cigarette users each purchasing once per month, yielding 66 million annual occasions for the all-flavored tax. The FDA-only upper bound uses 10.6 million occasions, reflecting current retail estimates that 16% of the flavored e-cigarette market is FDA-authorized (5.5M × 16% × 12). The demand-adjusted column scales each upper bound by the ratio of post-tax to pre-tax purchase shares over the full 36-month simulation horizon at $\beta = 1$; the 36-month horizon captures progressive habit-stock erosion that continues to depress demand beyond the first year. For the FDA-only tax, the ratio uses the FDA-authorized purchase share in both numerator and denominator, so it captures both demand reduction and substitution toward untaxed unauthorized alternatives.

Revenue is non-trivial in the aggregate. On a demand-adjusted basis at $\beta = 1$, the all-flavored tax generates substantially more revenue than the FDA-only tax: at the \$0.50/mL rate, \$98 million versus \$20 million per year. The gap is large because FDA-only taxation covers just 16% of the flavored market and induces substitution toward the

[54] For youth, the 2024 National Youth Tobacco Survey (NYTS) reports approximately 1.63 million current middle and high school e-cigarette users, of whom 87.6% use flavored products, yielding roughly 1.4 million flavored-using youth. For adults, the 2018 Tobacco Use Supplement to the Current Population Survey (TUS-CPS) reports a 1.6% current e-cigarette prevalence rate among U.S. adults, implying approximately 4 million adult vapers. Applying flavor-use shares of 89.8% among never-smoking vapers and 72.9% among current-smoker vapers yields approximately 4 million adult flavored e-cigarette users. Together these two populations sum to approximately 5.5 million.

[55] For example, at the \$0.50/mL all-flavored tax rate, the simulation implies that flavored purchase occasions fall to 18.6% of their pre-policy level on average over 36 months, yielding a demand-adjusted revenue of \$525M × 0.186 ≈ \$98M. For the FDA-only tax at the same rate, the ratio is 22.6%, applied to a base of 10.6 million FDA-authorized annual occasions, yielding \$89M × 0.226 ≈ \$20M.

untaxed unauthorized segment, compressing the taxable base from two directions. The ban generates zero revenue, so even the FDA-only \$0.50/mL tax compares favorably on this dimension. The revenue advantage of the all-flavored tax is a direct function of the unauthorized market's size, highlighting that aggregate revenue outcomes depend critically on how the tax is scoped relative to the unauthorized channel.

The preceding analysis focuses on intensive-margin effects among existing users. A distinct policy-relevant margin concerns new initiation: among households with essentially no prior flavored purchase history at the start of the counterfactual simulation, what share would make at least one flavored purchase within 36 months under the status quo, and how much does the tax reduce that probability?[56] Figure 19 plots the cumulative initiation rate, defined as the fraction of baseline non-users who have made at least one flavored purchase by each month, under the status quo and the \$0.50/mL tax, separately for TYA and non-TYA households at $\beta = 1$. To fix ideas, consider TYA households at $\beta = 1$: by month 6, 5.3% of baseline non-users have made their first flavored purchase under the status quo. The \$0.50/mL tax cuts this to 2.3%, more than halving the 6-month initiation probability. By month 36, the gap widens: 13.4% under the status quo versus 8.9% under the tax. Table 8 reports the full set of initiation rates at month 36 across policies and β values.

[56] Baseline non-users are defined as households whose simulated flavored habit stock satisfies $a^{\text{flav}} < 0.10$ at the start of the counterfactual. $a^{\text{flav}} \in [0, 1]$ is the standardized flavored habit stock used during estimation, constructed by simulating each household's purchase history forward through the observed data: it equals zero for a household that has never purchased flavored products and approaches one for a household with a long, uninterrupted flavored purchase history. The threshold 0.10 is chosen to be conservative: given the habit stock dynamics, a value below 0.10 implies the household has made at most one recent flavored purchase (if any), and its subsequent behavior is effectively that of a new initiator.

Figure 19. Cumulative Flavored E-Cigarette Initiation Rate Among Baseline Non-Users by TYA Status and Policy.

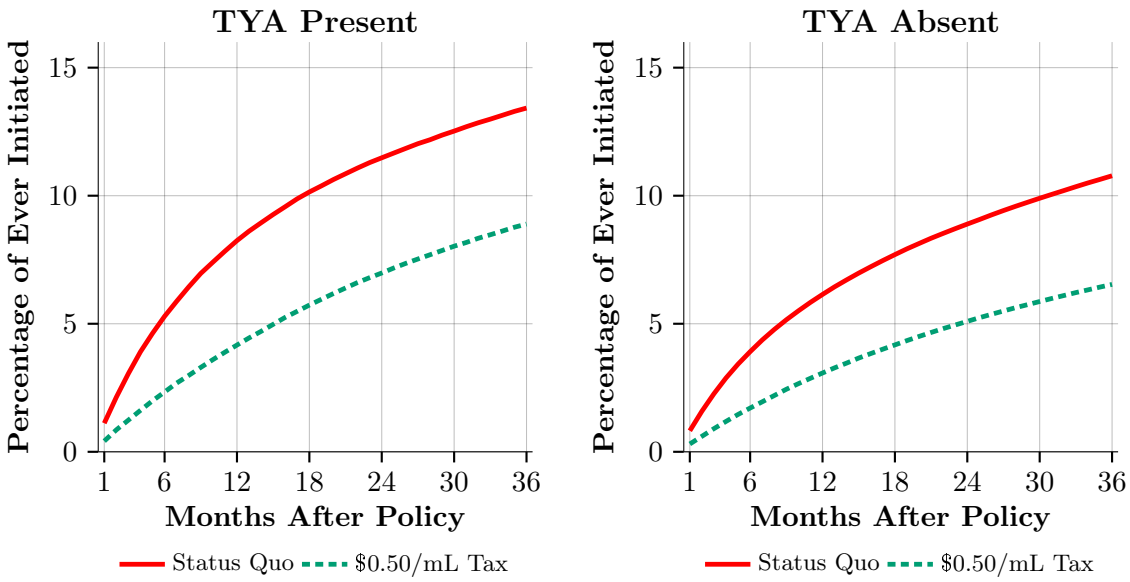

VIII.4. *Comparing Policy Instruments*

The results across the three policy instruments admit a clear welfare ranking. The comprehensive ban reduces addiction most aggressively but imposes the largest welfare cost and generates no revenue. The FDA-only ban limits welfare costs by targeting the smaller authorized market, but also limits effectiveness because consumers substitute freely to unauthorized alternatives. The \$0.50/mL all-flavored excise tax achieves nearly the same addiction reduction as the comprehensive ban while imposing substantially lower welfare costs, generating public revenue, and preserving the consumer's choice set. This section quantifies the efficiency advantage of the tax and documents how its benefits are distributed across consumer types.

The ban reduces addiction more than the tax at every horizon, and its advantage grows over time: under the ban, flavored products are unavailable and habit stocks erode continuously, whereas under the tax some consumers continue purchasing at the higher price, partially replenishing their habit stocks. Yet the welfare comparison strongly favors the tax. At $\beta = 1$, the \$0.50/mL tax achieves 92% of the ban's addiction reduction for TYA households (14.0% vs. 15.2%) at less than half the welfare cost (2.4% vs. 4.9%). This efficiency gap widens under present bias: at $\beta = 0.7$, the tax delivers 96% of the ban's addiction reduction at only 33% of its welfare cost. The intuition is that present-biased consumers systematically over-consume flavored products relative to their own long-run preferences, so a higher purchase price generates a sharper demand reduction among these households, part of which they experience as an internality correction.

Table 6 documents these patterns: addiction reduction is nearly flat across β while welfare costs fall at lower β as internality corrections grow, amplifying the tax's efficiency advantage (the TYA efficiency ratio rises from 5.78 at $\beta = 1$ to 11.53 at $\beta = 0.7$ under the tax, versus 3.13 to 3.97 under the ban).

Table 6. Addiction Reduction, Welfare Cost, and Efficiency Ratio at Period 36 Across β

	TYA Present			**TYA Absent**		
β	Addiction	Welfare	Efficiency	Addiction	Welfare	Efficiency
Comprehensive Flavor Ban						
0.7	−15.4%	−3.9%	3.97	−11.2%	−3.0%	3.68
1.0	−15.2%	−4.9%	3.13	−11.5%	−3.5%	3.23
$0.50/mL Excise Tax						
0.7	−14.9%	−1.3%	11.53	−10.8%	−1.1%	9.67
1.0	−14.0%	−2.4%	5.78	−10.6%	−1.8%	6.03

Entries report period-36 percentage changes in mean addiction and welfare relative to the same-β status quo. Efficiency is defined as |addiction reduction| / |welfare cost|: percentage points of addiction reduced per percentage point of welfare cost. Higher efficiency means more addiction reduction per unit of welfare sacrifice.

Table 7 summarizes welfare costs as percentage changes at periods 1 and 36 under each policy, by TYA status and $\beta \in \{0.7, 1.0\}$. Four features stand out. First, the comprehensive ban imposes the largest welfare cost at every specification: 7.3% for TYA households at $\beta = 1$ and period 1, compared with 4.0% for the FDA-only ban, 4.6% for the \$0.50/mL all-flavored tax, and 2.6% for the \$0.50/mL FDA-only tax. The FDA-only tax is least costly at every rate because consumers substitute freely to untaxed unauthorized alternatives. Second, TYA households bear higher welfare costs than non-TYA households under all policies reflecting a large flavor premium. Third, welfare costs fall substantially from period 1 to period 36 as habit stocks adjust; the decline is steeper under the ban (7.3% to 4.9% for TYA at $\beta = 1$) than under the tax because complete removal of flavored options accelerates habit-stock decay. Fourth, the $\beta = 0.7$ and $\beta = 1$ welfare costs are similar at period 1 but diverge at period 36: under the ban, time-consistent consumers ($\beta = 1$) bear a larger long-run welfare burden (4.9% versus 3.9% for TYA) because present-biased consumers place less weight on the future utility losses from habit-stock erosion, whereas under the tax at period 1, present-biased consumers bear higher immediate costs (4.9% versus 4.6%) because they continue purchasing at elevated prices.

Table 7. Percentage Change in Welfare at Periods 1 and 36, by Policy and β

		TYA Present		**TYA Absent**	
Policy	**Period**	$\beta = 0.7$	$\beta = 1.0$	$\beta = 0.7$	$\beta = 1.0$
Bans					
Comprehensive Ban	1	−7.4%	−7.3%	−5.8%	−5.7%
	36	−3.9%	−4.9%	−3.0%	−3.5%
FDA-Only Ban	1	−4.6%	−4.0%	−4.1%	−3.6%
	36	−2.0%	−2.1%	−1.8%	−1.9%
All-Flavored Tax					
\$0.10/mL	1	−2.9%	−2.2%	−2.3%	−1.8%
	36	−0.4%	−0.6%	−0.4%	−0.4%
\$0.25/mL	1	−3.9%	−3.5%	−3.1%	−2.7%
	36	−0.7%	−1.5%	−0.6%	−1.1%
\$0.50/mL	1	−4.9%	−4.6%	−3.9%	−3.7%
	36	−1.3%	−2.4%	−1.1%	−1.8%
FDA-Only Flavor Tax					
\$0.10/mL	1	−2.1%	−1.3%	−1.8%	−1.2%
	36	−0.9%	−0.2%	−0.7%	−0.2%
\$0.25/mL	1	−2.7%	−2.0%	−2.3%	−1.8%
	36	−0.9%	−0.5%	−0.7%	−0.5%
\$0.50/mL	1	−3.3%	−2.6%	−2.9%	−2.4%
	36	−1.0%	−1.0%	−0.9%	−0.9%

Entries report the percentage change in mean welfare relative to the same-β status quo. Welfare is the period-specific mean logsumexp value across simulated households and Monte Carlo draws. Period 1 is the first month of implementation; period 36 is the long-run (3-year) outcome.

Figure 20 traces the welfare comparison over time, confirming that the tax-ban efficiency advantage is persistent across all 36 simulated periods and both TYA subgroups.

Figure 20. Percentage Change in Welfare: \$0.50/mL Tax vs. Comprehensive Ban, by TYA Status and β.

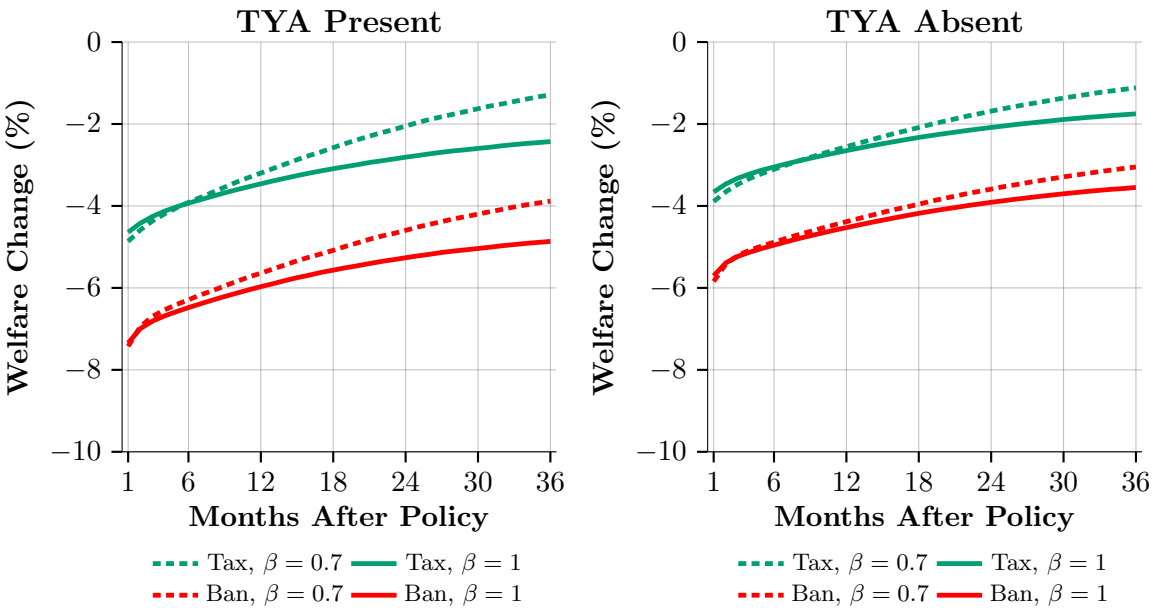

Table 8 extends the comparison to the extensive margin, reporting cumulative initiation rates among baseline non-users at month 36 across all policies and β values. Four findings stand out. First, TYA households have a higher baseline initiation rate: 13.4% versus 10.8% at $\beta = 1$, rising to 18.4% versus 15.9% at $\beta = 0.7$, consistent with their estimated flavor premium compounded by present-bias discounting of future addiction costs. Second, the \$0.50/mL tax outperforms the FDA-only ban as an initiation barrier for TYA households: at $\beta = 1$, the tax reduces TYA initiation from 13.4% to 8.9% (33.8% prevented) versus from 13.4% to 10.3% under the FDA-only ban (23.1% prevented); at $\beta = 0.7$, the corresponding figures are 38.6% versus 22.1% prevented. TYA households disproportionately purchase unauthorized brands that the FDA ban leaves untouched but the all-flavored tax prices up. Third, the tax is a more effective initiation deterrent under present bias: present-biased non-users facing a higher price are more likely to forgo a purchase they would have over-valued relative to their long-run preferences. Fourth, even the highest tax rate leaves roughly two-thirds of would-be TYA initiators unaffected at $\beta = 1$: the \$0.50/mL tax reduces the TYA initiation rate from 13.4% to 8.9%, preventing $(13.4 - 8.9)/13.4 \approx 34\%$ of initiations and leaving the remaining 66% unaffected. The comprehensive ban remains the only instrument that fully closes the initiation pathway.

Table 8. Cumulative Initiation Rate Among Baseline Non-Users at Month 36, by Policy and β

	TYA Present		TYA Absent	
Policy	$\beta = 0.7$	$\beta = 1.0$	$\beta = 0.7$	$\beta = 1.0$
Status Quo	18.36%	13.42%	15.88%	10.78%
FDA-Only Ban	14.30%	10.32%	10.49%	6.92%
\$0.10/mL Tax	16.20%	12.27%	13.56%	9.56%
\$0.25/mL Tax	13.92%	10.81%	11.28%	8.24%
\$0.50/mL Tax	11.28%	8.89%	8.71%	6.54%
Comprehensive Ban	0%	0%	0%	0%

Entries report the fraction of baseline non-user households who make at least one flavored e-cigarette purchase by month 36, averaged across 100 Monte Carlo draws per household.

The aggregate addiction and welfare figures mask substantial heterogeneity across the two latent consumer types identified in estimation. Table 9 reports addiction and welfare changes at periods 1, 12, and 36 for Type 1 and Type 2 households under the comprehensive ban and the \$0.50/mL tax at $\beta = 1$.

Table 9. Addiction and Welfare Changes by Consumer Type: Comprehensive Ban and \$0.50/mL Tax, $\beta = 1$

	Period 1		Period 12		Period 36	
	Addiction	Welfare	Addiction	Welfare	Addiction	Welfare
Comprehensive Flavor Ban						
Type 1 ($\approx 90\%$)	-0.5%	-1.2%	-1.5%	-1.1%	-3.1%	-0.9%
Type 2 ($\approx 10\%$)	-29.3%	-64.7%	-60.9%	-50.2%	-84.4%	-39.2%
\$0.50/mL Excise Tax						
Type 2 ($\approx 10\%$)	-27.0%	-41.4%	-58.8%	-28.8%	-80.9%	-18.5%

Entries report the percentage change in mean addiction and welfare at each period relative to the same-β status quo, evaluated at $\beta = 1$. Type 1 comprises approximately 90% of the sample and represents households with strong baseline cigarette preferences. Type 2 ($\approx$10%) represents active e-cigarette vapers with high flavor preferences, for whom flavor restrictions bind most directly. Type 1 results under the tax are omitted because the tax has negligible effects on cigarette-oriented households.

The heterogeneity is stark. Type 1 households (strong cigarette preferences, low e-cigarette base utility) experience only a 3.1% addiction reduction and less than 1% welfare loss at period 36: removing flavored e-cigarettes barely perturbs their equilibrium. Type 2 households experience an 84.4% addiction reduction and a 39.2% welfare loss, since flavored e-cigarettes are central to their consumption and the flavored habit

stock erodes almost entirely over 36 months.

Within Type 2, the addiction difference between the ban and the \$0.50/mL tax is small at period 36 (84.4% vs. 80.9%), but the welfare difference is large (39.2% vs. 18.5%): the tax preserves option value for consumers who choose to pay. The aggregate efficiency advantage of the tax accrues almost entirely to Type 2 households, precisely those most affected by the restriction.

IX. Conclusion

This paper develops and estimates a dynamic discrete choice model of cigarette and e-cigarette demand to evaluate three active e-cigarette flavor policies: a comprehensive ban on all flavored products, a ban limited to FDA-authorized brands, and a per-milliliter excise tax. The counterfactual simulations incorporate both nicotine addiction stocks and a flavored habit stock that governs how quickly displaced consumers exit the flavored market.

Three findings emerge. First, the feared backfire from a flavor ban does not occur. The dominant behavioral response is market exit, not cigarette substitution: the outside-option share rises by 3.9 percentage points at 36 months under the comprehensive ban, while cigarette shares fall modestly and return to below-baseline levels by the final period. Without flavored options to sustain the habit stock, habit-dependent demand for nicotine products decays and market exit becomes progressively more attractive.

Second, enforcement scope is decisive. The comprehensive ban reduces mean nicotine addiction by 12.2% over 36 months, while the FDA-only ban captures only 84% of that impact, with the foregone reduction driven entirely by substitution to unauthorized flavored products that remain available. The attenuation is sharpest for TYA households, for whom the FDA-only ban captures only 82% of the comprehensive ban's addiction reduction. These estimates likely understate the true enforcement gap: in the 2021–2023 estimation sample, approximately 58–65% of flavored purchases flowed through unauthorized channels, but national retail data indicate that share has since risen to 86% as of 2024. As the unauthorized market grows, the gap between a comprehensive and an authorized-only ban widens further, making the estimates here a lower bound on what a realistic FDA-only ban would achieve today.

Third, a \$0.50/mL excise tax achieves 92% of the comprehensive ban's addiction reduction at roughly half the welfare cost and generates an estimated \$525 million annually in public revenue. Present bias does not reverse this ranking. When consumers are present-biased, they over-consume flavored e-cigarettes relative to their own long-run preferences, so part of the measured welfare loss from either policy reflects an internality correction rather than genuine forgone utility. This lowers the welfare cost of both instruments at $\beta = 0.7$ relative to $\beta = 1$, but it does so proportionally more for the tax, widening its efficiency advantage under stronger present bias.

Limitations. Several features of the data and model bound the scope of these

conclusions. First, the HMS data are collected at the household level, so TYA status captures the presence of a person aged 13–25 in the household rather than confirming who is purchasing. The TYA estimates likely understate true teen demand since informal and online transactions are likely underestimated in retail scanner data. Second, the model excludes nicotine pouches (Zyn, on!, Velo), which are available in a wide range of flavors and represent a plausible substitution margin following a flavor ban on e-cigarettes; to the extent pouches are close substitutes, the addiction reduction and market exit estimates are upper bounds. Third, the quasi-hyperbolic discount factor β is calibrated rather than identified from the data, so the present-bias analysis is best understood as a welfare sensitivity exercise rather than a structural estimate of consumer time preferences. Fourth, the supply side is static: equilibrium prices follow from a Bertrand–Nash first-order condition with fixed costs and offerings, abstracting from dynamic pricing, firm entry, or product repositioning that a ban or tax could trigger.

Real-world implications. The enforcement gap finding implies that flavor restrictions imposed through the PMTA process will achieve only a fraction of their intended impact as long as unauthorized products remain on retail shelves. Regulatory resources directed at expanding enforcement against unauthorized brands would likely yield a higher return in addiction reduction per dollar than additional restrictions on the small authorized segment. On the tax side, the Tobacco Tax Equity Act of 2023 would translate to roughly \$2.78 per milliliter for high-nicotine disposables; the results here show that even a \$0.50/mL tax, far below the proposed rate, achieves substantial addiction reduction at a welfare cost below two-thirds that of an outright ban while generating meaningful public revenue. Like a ban, however, a tax faces the same evasion pressure from unauthorized firms: because unauthorized manufacturers already operate outside the regulatory system, they can simply avoid remitting the tax, leaving the price of unauthorized alternatives largely unchanged and preserving the competitive advantage of non-compliant firms.

Implications for future research. Three extensions follow naturally. Incorporating nicotine pouches into the demand system would test whether a flavor-specific e-cigarette ban simply redirects consumption to an unregulated oral nicotine market. Identifying β from the data, rather than calibrating it, would require exogenous variation in the long-run cost of nicotine dependence, such as plausibly exogenous state cigarette tax changes, and would transform the present-bias analysis from a sensitivity check into an integral part of the structural model. Finally, a dynamic supply side allowing for firm entry, product repositioning, and strategic pricing responses would refine the enforcement gap estimates by modeling how the unauthorized market segment expands in response to partial bans.

References

Adda, **Jérôme** and **Francesca Cornaglia** (2006). "Taxes, Cigarette Consumption, and Smoking Intensity". *The American Economic Review* 96.4, pp. 1013–1028. DOI: https://doi.org/10.1257/aer.96.4.1013.

Aguirregabiria, **Victor** and **Pedro Mira** (2010). "Dynamic Discrete Choice Structural Models: A Survey". *Journal of Econometrics* 156, pp. 38–67. DOI: http://dx.doi.org/10.1016/j.jeconom.2009.09.007.

Babb, **Stephen**, **Ann Malarcher**, **Gillian Schauer**, **Kat Asman**, and **Ahmed Jamal** (2017). "Quitting Smoking Among Adults — United States, 2000–2015". *MMWR Morbidity and Mortality Weekly Report* 65.52, pp. 1457–1464. DOI: https://doi.org/10.15585/mmwr.mm6552a1.

Baltagi, **Badi H.** and **James M. Griffin** (2002). "Rational Addiction to Alcohol: Panel Data Analysis of Liquor Consumption". *Health Economics* 11.6, pp. 485–491. DOI: https://doi.org/10.1002/hec.748.

Bask, **Mikael** and **Maria Melkersson** (2004). "Rationally Addicted to Drinking and Smoking?" *Applied Economics* 36.4, pp. 373–381. DOI: https://doi.org/10.1080/00036840410001674295.

Becker, **Gary S.**, **Michael Grossman**, and **Kevin M. Murphy** (1994). "An Empirical Analysis of Cigarette Addiction". *The American Economic Review* 84.3, pp. 396–418. URL: https://www.jstor.org/stable/2118059.

Becker, **Gary S.** and **Kevin M. Murphy** (1988). "A Theory of Rational Addiction". *Journal of Political Economy* 96.4, pp. 675–700. DOI: https://doi.org/10.1086/261558.

Benowitz, **Neal L.**, **Janne Hukkanen**, and **Peyton Jacob** (2009). "Nicotine Chemistry, Metabolism, Kinetics, and Biomarkers". *Nicotine Psychopharmacology* 192, pp. 29–60. DOI: https://doi.org/10.1007/978-3-540-69248-5_2.

Berry, **Steven**, **James Levinsohn**, and **Ariel Pakes** (1995). "Automobile Prices in Market Equilibrium". *Econometrica* 63.4, pp. 841–890. DOI: https://doi.org/10.2307/2171802.

Chaloupka, **Frank** (1991). "Rational Addictive Behavior and Cigarette Smoking". *Journal of Political Economy* 99.4, pp. 722–742. URL: https://www.jstor.org/stable/2937778.

Chen, **Jiawei** and **Vithala R. Rao** (2020). "A Dynamic Model of Rational Addiction with Stockpiling and Learning: An Empirical Examination of E-cigarettes". *Management Science* 66.12, pp. 5886–5905. DOI: https://doi.org/10.1287/mnsc.2019.3490.

Chen, **Jiawei** and **Colin Reinhardt** (2025). "Flavorants and Addiction: An Empirical Analysis of Cigarette Bans and Taxation". *International Journal of Industrial Organization* 99.103147. DOI: https://doi.org/10.1016/j.ijindorg.2025.103147.

Cosgrove, **Kelly P.**, **Jeffery Batis**, **Frederic Bois**, **Paul K. Maciejewski**, **Irina Esterlis**, **Tracy Kloczynski**, **Stephanie Stiklus**, **Suchitra Krishnan-Sarin**, **Stephanie O'Malley**, **Edward Perry**, **Gilles Tamagnan**, **John P. Seibyl**, and **Julie K. Staley** (2009). "β_2-Nicotinic Acetylcholine Receptor Availability during Acute and Prolonged Abstinence from Tobacco Smoking". *Archives of General Psychiatry* 66.6, pp. 666–676. DOI: https://doi.org/10.1001/archgenpsychiatry.2009.41.

Cotti, **Chad D.**, **Charles J. Courtemanche**, **Yang Liang**, **Johanna Catherine Maclean**, **Erik T. Nesson**, and **Joseph J. Sabia** (2025). "The Effect of E-Cigarette Flavor Bans on Tobacco Use". *Journal of Health Economics* 102.103013. DOI: https://doi.org/10.1016/j.jhealeco.2025.103013.

Deng, **Xueting**, **Yuqing Zheng**, and **J.S. Butler** (2023). "An Empirical Analysis of E-Cigarette Addiction". *Journal of Applied Economics* 26.1. DOI: https://doi.org/10.1080/15140326.2023.2223953.

Erdem, **Tülin**, **Susumu Imai**, and **Michael P. Keane** (2003). "Brand and Quantity Choice Dynamics Under Price Uncertainty". *Quantitative Marketing and Economics* 1.1, pp. 103–124. DOI: https://doi.org/10.1257/aer.89.1.103.

Friedman, **Abigail S.**, **Michael F. Pesko**, and **Travis R. Whitacre** (2024). "Flavored E-Cigarette Sales Restrictions and Young Adult Tobacco Use". *JAMA Health Forum* 5.12. DOI: https://doi.org/10.1001/jamahealthforum.2024.4594.

Gentzkow, **Matthew** (2007). "Valuing New Goods in a Model with Complementarity: Online Newspapers". *The American Economic Review* 97.3, pp. 713–744. DOI: https://doi.org/10.1257/aer.97.3.713.

Gordon, **Brett R.** and **Baohong Sun** (2015). "A Dynamic Model of Rational Addiction: Evaluating Cigarette Taxes". *Marketing Science* 34.3, pp. 452–460. DOI: https://doi.org/10.1287/mksc.2014.0885.

Gowrisankaran, **Gautam** and **Marc Rysman** (2020). "A Framework for Empirical Models of Dynamic Demand". URL: https://drive.google.com/file/d/1jYrpJoWfNgofsi39Nv4bYFBKDj1_Nwxg/view.

— (2012). "Dynamics of Consumer Demand for New Durable Goods". *Journal of Political Economy* 120.6, pp. 1173–1219. DOI: https://doi.org/10.1086/669540.

Gruber, **John** and **Botond Köszegi** (2001). "Is Addiction "Rational"? Theory and Evidence". *The Quarterly Journal of Economics* 116.4, pp. 1261–1303. DOI: https://doi.org/10.1162/003355301753265570.

— (2004). "Tax Incidence When Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes". *Journal of Public Economics* 88.9-10, pp. 1959–1987. DOI: https://doi.org/10.1016/j.jpubeco.2003.06.001.

Hendel, **Igal** and **Aviv Nevo** (2006). "Measuring the Implications of Sales and Consumer Inventory Behavior". *Econometrica* 74.6, pp. 1637–1673. DOI: https://doi.org/10.1111/j.1468-0262.2006.00721.x.

Hortaçsu, **Ali** and **Joonhwi Joo** (2023). *Structural Econometric Modeling in Industrial Organization and Quantitative Marketing: Theory and Applications*. DOI: https://doi.org/10.2307/jj.5425946.

Hughes, **John R.**, **Josue Keely**, and **Shelly Naud** (2004). "Shape of the Relapse Curve and Long-Term Abstinence Among Untreated Smokers". *Addiction* 99.1, pp. 29–38. DOI: https://doi.org/10.1111/j.1360-0443.2004.00540.x.

Hughes, **John R.**, **Erica N. Peters**, and **Shelly Naud** (2008). "Relapse to Smoking After 1 Year of Abstinence: A Meta-Analysis". *Addictive Behaviors* 33.12, pp. 1516–1520. DOI: https://doi.org/10.1016/j.addbeh.2008.05.012.

Kechter, **Afton**, **Jessica Cho**, **Richard A. Miech**, **Jessica L. Barrington-Trimis**, and **Adam M. Leventhal** (2021). "Nicotine Dependence Symptoms in U.S. Youth Who Use JUUL E-Cigarettes". *Drug and Alcohol Dependence* 227, p. 108941. DOI: https://doi.org/10.1016/j.drugalcdep.2021.108941.

Krall, **Elizabeth A.**, **Arthur J. Garvey**, and **Raul I. Garcia** (2002). "Smoking Relapse After 2 Years of Abstinence: Findings from the VA Normative Aging Study". *Nicotine and Tobacco Research* 4.1, pp. 95–100. DOI: https://doi.org/10.1080/14622200110098428.

Laibson, **David** (1997). "Golden Eggs and Hyperbolic Discounting". *The Quarterly Journal of Economics* 112.2, pp. 443–477. DOI: https://doi.org/10.1162/003355397555253.

Mamede, **Marcos**, **Keiichi Ishizu**, **Masashi Ueda**, **Takashi Mukai**, **Yoshiharu Iida**, **Hidehiko Kawashima**, **Hideo Fukuyama**, **Kaori Togashi**, and **Tsunehiko Nishimura** (2007). "Temporal Change in Human Nicotinic Acetylcholine Receptor After Smoking Cessation: 5IA SPECT Study". *Journal of Nuclear Medicine* 48.11, pp. 1829–1835. DOI: https://doi.org/10.2967/jnumed.107.043471.

McFadden, **Daniel** (1972). *Conditional Logit Analysis of Qualitative Choice Behavior*. URL: https://eml.berkeley.edu/reprints/mcfadden/zarembka.pdf.

McLaughlin, **Ian**, **John A. Dani**, and **Mariella De Biasi** (2015). "Nicotine Withdrawal". In: *The Neuropharmacology of Nicotine Dependence*. Vol. 24. Current Topics in Behavioral Neurosciences. Springer, pp. 99–123. DOI: https://doi.org/10.1007/978-3-319-13482-6_4.

O'Donoghue, **Ted** and **Matthew Rabin** (1999). "Doing It Now or Later". *The American Economic Review* 89.1, pp. 103–124. DOI: https://doi.org/10.1257/aer.89.1.103.

Petrin, **Amil** (2002). "Quantifying the Benefits of New Products: The Case of the Minivan". *Journal of Political Economy* 110.4, pp. 705–729. DOI: https://doi.org/10.1086/340779.

Piccoli, **Luca** and **Silvia Tiezzi** (2021). "Rational Addiction and Time-Consistency: An Empirical Test". *Journal of Health Economics* 80.102546. DOI: https://doi.org/10.1016/j.jhealeco.2021.102546.

Press, **William H.**, **Saul A. Teukolsky**, **William T. Vetterling**, and **Brian P. Flannery** (2007). *Numerical Recipes: The Art of Scientific Computing*.

Prochaska, **Judith L.**, **Erin A. Vogel**, and **Neal L. Benowitz** (2022). "Nicotine Delivery and Cigarette Equivalents from Vaping a JUUL Pod". *Tobacco Control* 31, pp. 88–93. DOI: https://doi.org/10.1136/tobaccocontrol-2020-056367.

Rust, **John** (1987). "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher". *Econometrica* 55.5, pp. 999–103. URL: http://www.jstor.org/stable/1911259?origin=JSTOR-pdf.

Saffer, **Henry**, **Selen Ozdogan**, **Michael Grossman**, **Daniel L. Dench**, and **Dhaval M. Dave** (2025). "Comprehensive E-Cigarette Flavor Bans and Tobacco Use Among Youth, Young Adults, and Adults". *National Bureau of Economics Research (NBER)* 32534. DOI: https://doi.org/10.3386/w32534.

Stigler, **George J.** and **Gary S. Becker** (1977). "De Gustibus Non Est Disputandum". *The American Economic Review* 67.2, pp. 76–90. URL: https://www.jstor.org/stable/1807222.

Tsai, **James**, **Kimp Walton**, **Blair N. Coleman**, **Saida R. Sharapova**, **Sarah E. Johnson**, **Sara M. Kennedy**, and **Ralph S. Caraballo** (2016). "Reasons for Electronic Cigarette Use Among Middle and High School Students". *Morbidity and Mortality Weekly Report* 67.6, pp. 196–200. DOI: https://doi.org/10.15585/mmwr.mm6706a5.

Yingst, **Jessica M.**, **Shari Hrabovsky**, **Andrea Hobkirk**, **Neil Trushin**, **John P. Richie Jr.**, and **Jonathan Foulds** (2019). "Nicotine Absorption Profile Among Regular Users of a Pod-Based Electronic Nicotine Delivery System". *JAMA Network Open* 2.11. DOI: https://doi.org/10.1001/jamanetworkopen.2019.15494.

X. Appendix

X.1. *Figures*

Figure A1. Post-Purchase Paths After a Cigarette Purchase.

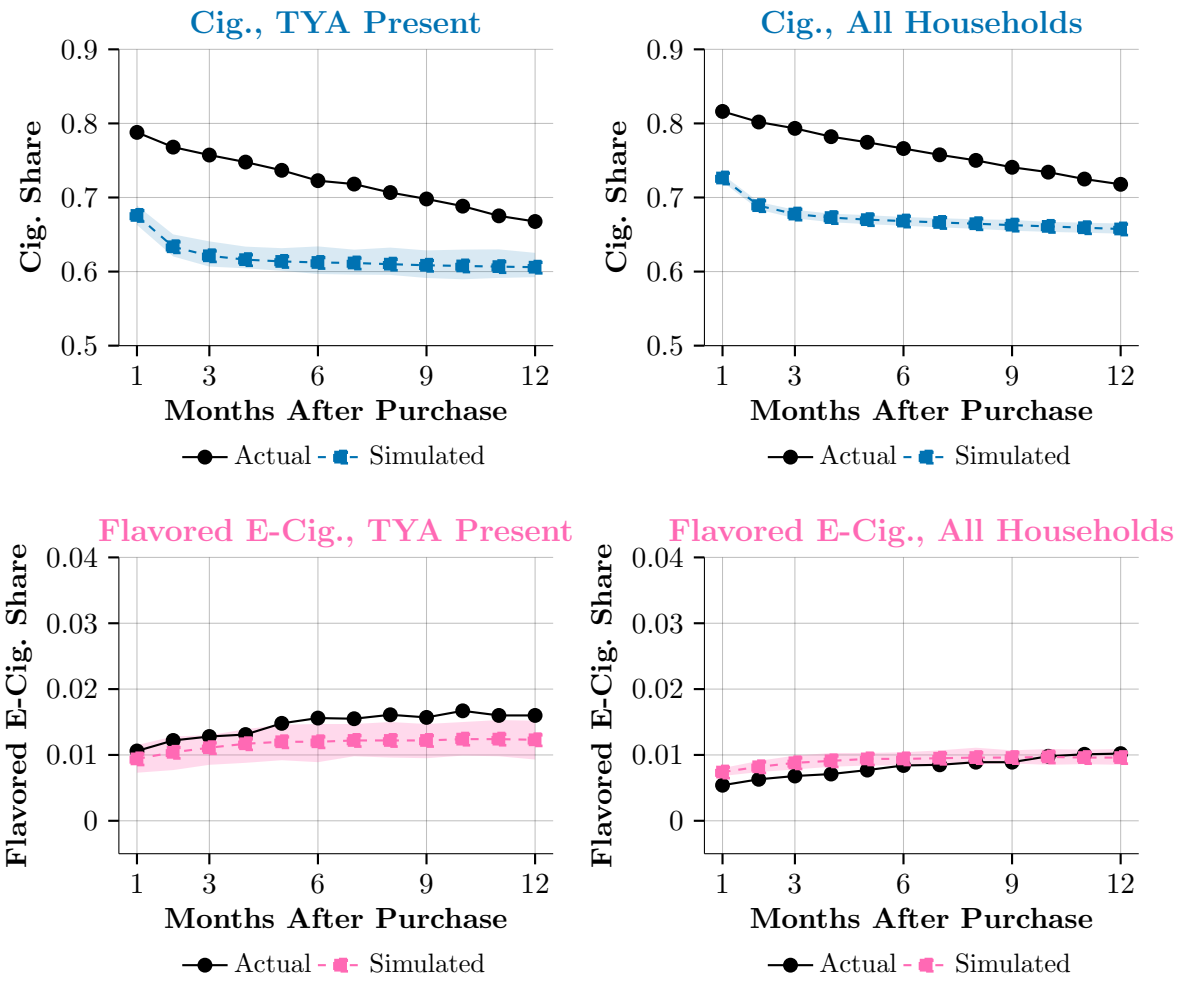

This figure plots destination shares at each month following an initial cigarette purchase, separately for TYA-present households (left column) and all households (right column). The top row shows cigarette shares and the bottom row shows flavored e-cigarette shares. Black circles show actual shares computed from observed choices. Colored dashed lines with squares show simulated means across $S = 100$ forward simulation draws in which the addiction stock evolves endogenously based on simulated choices. Shaded bands show 95% simulation confidence intervals.

Figure A1 performs the post-purchase path analysis starting from a cigarette purchase, plotting destination shares at each horizon for cigarettes and flavored e-cigarettes separately by TYA status. Actual cigarette shares at horizon 1 are 78.8% for TYA households and 81.6% overall, declining to 66.8% and 71.8% by month 12. Simulated shares begin somewhat lower at 67.6% (TYA) and 72.6% (overall), reflecting the model's starting-point underprediction of cigarette lock-in, and converge gradually toward the actual trajectory as the addiction stock builds along simulated paths. Flows toward e-

cigarettes and flavored e-cigarettes following a cigarette purchase are near zero in both the data and the model throughout, confirming that cross-category substitution from cigarettes to e-cigarettes is rare over the medium run. The figure shows cigarette persistence is consistently underpredicted by roughly ten percentage points, which likely occurs given there is not a cigarette habit stock analogous to the flavor habit stock already present in the model. However, the model does capture dynamic substitution patterns from cigarettes to e-cigarettes quite well, despite the small fraction of individuals making this switch. The latter is more important for establishing how the model can evaluate how e-cigarette flavor policy impacts the extensive and intensive behavioral margins.

Figure A2. Share of Purchases by Product Type.

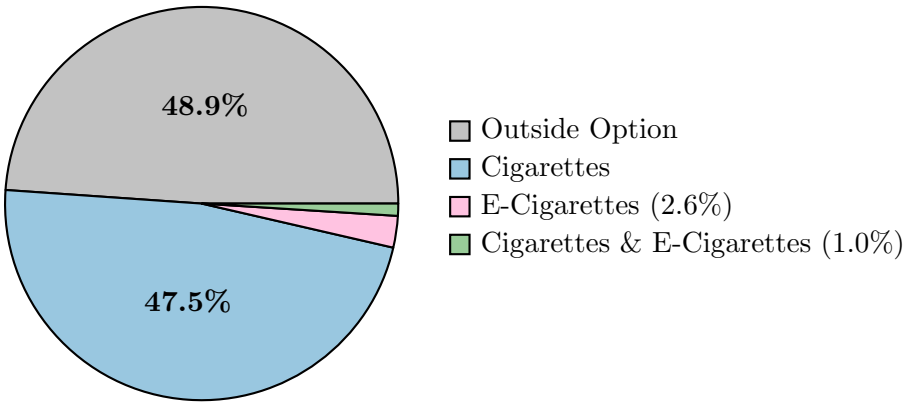

Figure A3. Inflation-Adjusted Median Monthly Prices of Packs of Cigarettes (Panel (A)) and Milliliters of E-Liquid (Panel (B)).

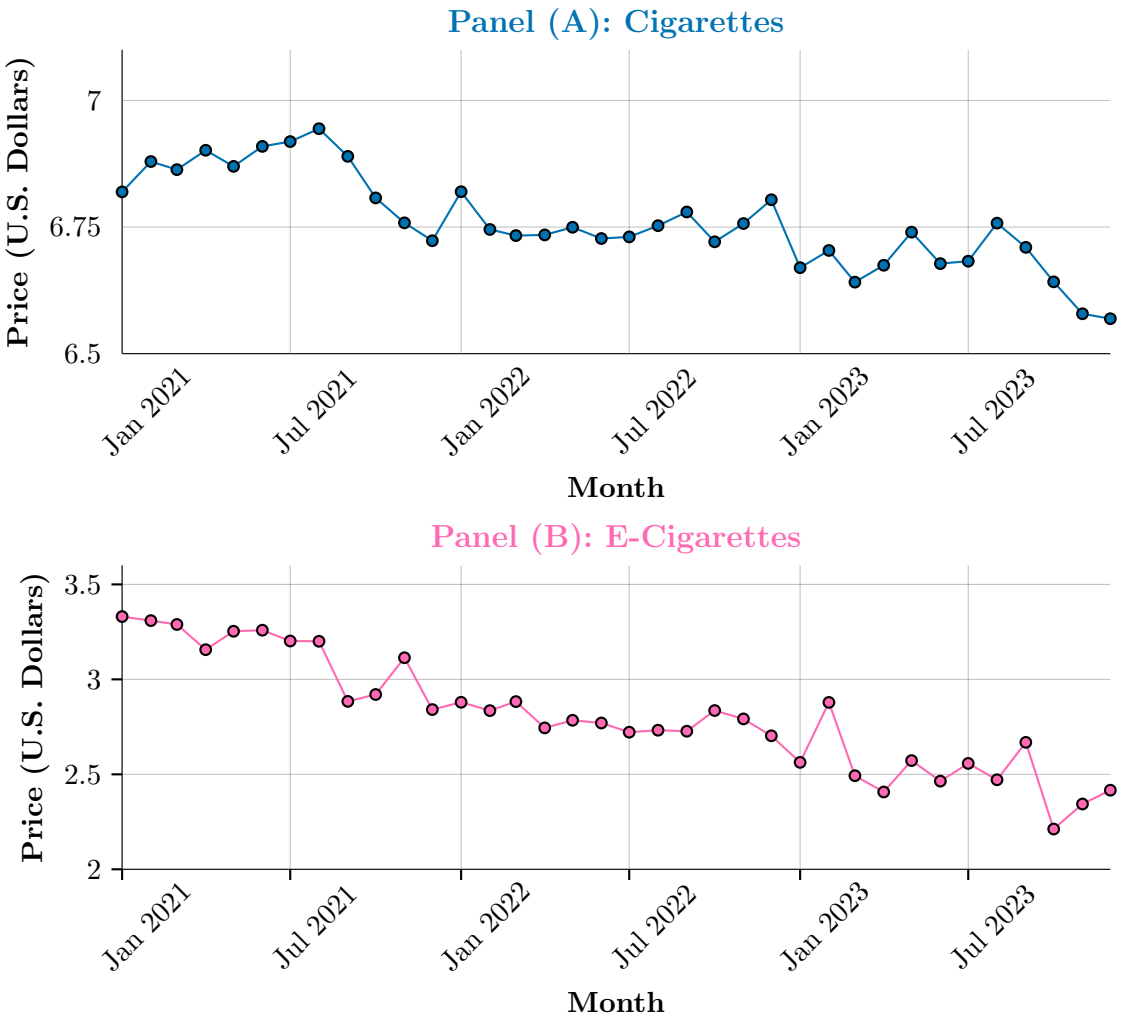

Inflation-adjusted prices use the consumer price index (CPI) for tobacco products found at https://fred.stlouisfed.org/series/CUSR0000SEGA.

X.2. *Tables*

X.2.1. *Variable Descriptions*

Table A.1. Data Description for Selected Variables.

Variable	Description
Household Income	Annual household income in U.S. dollars, top-coded at $100,000.
HoH Age	Mean head-of-household (HoH) age across household head(s).
HoH w/ College Degree	Indicator for whether at least one head-of-household (HoH) has a college degree.
White Households	Indicator for whether the household is composed of white individuals.
Married Households	Indicator for whether the household contains married individuals.
Children Present	Indicator for whether the household contains child member aged 0 – 12, inclusive.
Teen Present	Indicator for whether the household contains a teenage member aged 13 – 18, inclusive.
Young Adult Present	Indicator for whether household contains a young adult member aged 19 – 25, inclusive.
Son or Daughter Present	Indicator for whether the son or daughter of the household head(s) is present.
Inside Good/Option	Refers to cigarettes of any kind and e-cigarettes in cartridge, liquid, or disposable form.
Outside Good/Option	Refers to all possible non-inside goods.
Cig.	Refers to cigarettes.
E-Cig.	Refers to e-cigarettes.

X.2.2. *Teen and Young Adult Regressions*

Table A.2. OLS Results for the Effect of a Teen in the Household on Purchasing Flavored E-Cigarettes.

Panel (A): All Household-Months for Inside Good Households

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Teen	0.015***	0.016***	0.016***
	(0.004)	(0.005)	(0.005)
HH Size		−0.0004	−0.0005
		(0.001)	(0.001)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	201,348	201,348	201,348
R^2	0.002	0.002	0.005

Panel (B): All Household Months Containing an E-Cig. Purchase

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Teen	0.145***	0.118**	0.125**
	(0.042)	(0.048)	(0.049)
HH Size		0.016	0.008
		(0.014)	(0.014)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	7,207	7,207	7,207
R^2	0.020	0.021	0.122

Panel (A) includes all household-months for those households making at least one cigarette or e-cigarette purchase over the 2021 – 2023 time frame. Panel (B) restricts the sample to household-months containing e-cigarette purchases. Cluster-robust standard errors are at the household level and reported in parentheses. Significance levels: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.3. OLS Results for the Effect of a Young Adult in the Household on Purchasing Flavored E-Cigarettes.

Panel (A): All Household-Months for Inside Good Households

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Young Adult	0.013***	0.012**	0.012**
	(0.004)	(0.005)	(0.005)
HH Size		0.0005	0.0004
		(0.001)	(0.001)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	201,348	201,348	201,348
R^2	0.002	0.002	0.004

Panel (B): All Household Months Containing an E-Cig. Purchase

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Young Adult	0.121***	0.087*	0.102**
	(0.044)	(0.047)	(0.043)
HH Size		0.023*	0.015
		(0.013)	(0.013)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	7,207	7,207	7,207
R^2	0.015	0.018	0.120

Panel (A) includes all household-months for those households making at least one cigarette or e-cigarette purchase over the 2021 – 2023 time frame. Panel (B) restricts the sample to household-months containing e-cigarette purchases. Cluster-robust standard errors are at the household level and reported in parentheses. Significance levels: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.4. OLS Results for the Effect of Teens and Young Adults on Purchasing Flavored E-Cigarettes: All Household-Months.

Panel (A): HH Contains a Teen

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Teen	0.002***	0.002***	0.002***
	(0.001)	(0.001)	(0.001)
HH Size		0.00001	0.00002
		(0.0001)	(0.0001)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	1,985,724	1,985,724	1,985,724

Panel (B): HH Contains a Young Adult

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Young Adult	0.001***	0.001**	0.001**
	(0.001)	(0.001)	(0.001)
HH Size		0.0001	0.0001
		(0.0001)	(0.0001)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	1,985,724	1,985,724	1,985,724

Panel (C): HH Contains a Teen or Young Adult

	Purchased Flavored E-Cig.		
	(1)	(2)	(3)
HH Contains a Teen or Young Adult	0.002***	0.002***	0.002***
	(0.0004)	(0.001)	(0.001)
HH Size		−0.0001	−0.0001
		(0.0001)	(0.0001)
State Fixed Effects			✓
Month Fixed Effects	✓	✓	✓
Observations	1,985,724	1,985,724	1,985,724

All panels include all household-months, including non-purchasing households. Panel (A) uses an indicator for whether the household contains a teen (aged 13–18) as the key regressor. Panel (B) uses an indicator for a young adult (aged 19–25). Panel (C) uses an indicator for either a teen or young adult. Cluster-robust standard errors are at the household level and reported in parentheses. Significance levels: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

X.2.3. *Stockpiling Test*

A key identifying assumption in the structural model is that the persistence in household tobacco purchases reflects habit formation and addiction rather than stockpiling. A household that purchases a large quantity in one month has drawn down its future demand, making a zero-purchase month more likely in the next period. Addiction operates in precisely the opposite direction: a high-purchase month builds habit stock, raising the marginal utility of consumption and making a zero-purchase month less likely going forward. These opposing predictions yield a reduced-form test. Regressing an indicator for a zero-purchase month on lagged consumption quantity, controlling for household and month fixed effects, delivers a coefficient whose sign tests for the correct mechanism. Table A.5 reports these results for both cigarettes and e-cigarettes. In both panels and both specifications, the coefficient on lagged consumption is negative and significant, suggesting stockpiling is not a prevalent feature of the data and supporting the addiction interpretation incorporated in the structural model.

Table A.5. Zero-Purchase Transition Test for Stockpiling.

Panel (A): Cigarettes

	Zero Cigarette Purchase Month	
	(1)	(2)
Lagged Packs Purchased	-0.0078^{***}	-0.0022^{***}
	(0.0008)	(0.0002)
Household Fixed Effects		✓
Month Fixed Effects	✓	✓
Observations	191,412	191,412

Panel (B): E-Cigarettes

	Zero E-Cigarette Purchase Month	
	(1)	(2)
Lagged mL of E-liquid Purchased	-0.0052^{***}	-0.0011^{***}
	(0.0012)	(0.0004)
Household Fixed Effects		✓
Month Fixed Effects	✓	✓
Observations	191,412	191,412

The sample includes all household-months for inside good households. The dependent variable is an indicator equal to one if the household made zero purchases of the given product in month t. The key regressor is total consumption of that product in month $t-1$, set to missing when months $t-1$ and t are non-consecutive. Column (1) includes month fixed effects only; column (2) adds household fixed effects (two-way within estimator). Cluster-robust standard errors are at the household level and reported in parentheses. Significance levels: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

X.2.4. *Abstention Spell Analysis*

Table A.6. Abstention Spell Summary by Product

	Cig.	**E-Cig.**
At-Risk Months	94,352	6,903
Abstention Spells	21,388	2,461
Monthly Abstention Rate (%)	22.67	35.65
Resumption Rate (%)	72.53	59.81
Within 1 Month (%)	49.94	45.18
Within 3 Month (%)	78.51	74.66

An at-risk month is a household-month in which the household purchased the product in the previous month. An abstention spell occurs when an at-risk household does not purchase in the current month. Monthly abstention rate is the fraction of at-risk months that result in an abstention spell. Resumption rate is the fraction of abstention spells that end in a return to purchasing within the panel window. Within 1 month and within 3 months report the fraction of resumed spells in which the household returned to purchasing within 1 or 3 months of the abstention spell, respectively. Cigarette and e-cigarette spells each include months involving a cigarette-and-e-cigarette bundle.

X.2.5. *AR(1) Price Process Estimates*

Table A.7. Estimated Monthly Price Processes and Covariance Matrix.

Panel (A): AR(1) Price Process Estimates $\left(\widehat{\boldsymbol{\phi_k}}\right)$

	Cigarettes	E-Cigarettes
$\widehat{\phi}_0$	0.918	0.341
	(0.589)	(0.234)
$\widehat{\phi}_1$	0.864	0.880
	(0.087)	(0.083)
Log likelihood	55.66	15.00
AIC	−105.31	−24.00

Panel (B): Variance–Covariance Matrix $\left(\widehat{\Sigma}\right)$

	Cigarettes	E-Cigarettes
Cigarettes	0.0026	.
E-Cigarettes	0.0025	0.0247

Panel (A) reports autoregressive estimates for monthly market-level median real prices. Both cigarette and e-cigarette prices follow AR(1) processes in real per-pack and per-mL prices, respectively. Means for each process are 6.74 and 2.82, respectively. Standard errors are reported in parentheses, with the standard error for $\widehat{\phi}_0$ computed using the Delta Method. Panel (B) reports the variance–covariance matrix of the estimated AR innovations.

Estimated price coefficients display substantial heterogeneity across cigarette and e-cigarette products. Monthly cigarette price shocks exhibit very low volatility, with an estimated standard deviation of roughly 0.05, implying smooth price adjustments. In contrast, e-cigarette prices are considerably more volatile, with an estimated standard deviation of approximately 0.16, more than three times larger than that of cigarettes. Despite these differences, price shocks across the two markets are moderately correlated, with an estimated covariance of 0.00251 corresponding to a correlation of roughly 0.31. This positive comovement indicates the presence of common wholesale cost shocks, shared inflationary pressures, and macroeconomic influences affecting both products simultaneously. Overall, cigarette prices follow a highly stable and predictable process, whereas e-cigarette prices exhibit greater short-run uncertainty.

X.3. *Dynamic Programming Problem Details*

I outline the details of the dynamic programming problem in this section.[57] I consider an infinite time horizon $t \in \{1, 2, \dots N\}$.[58] In each time period t, the consumer $i \in \{1, 2, \dots, \}$ observes the state vector $\boldsymbol{s}_{it} = (\boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}) = \left(h_{it}, a^f_{it}, a^s_{it}, a^{\text{flav}}_{it}, \boldsymbol{p}_{it}, \boldsymbol{\varepsilon}_{it}\right)$, forms expectations of future prices, and chooses the action $\boldsymbol{d}_{it}$ to maximize the expected quasi-hyperbolic discounted sum of flow utilities defined as

$$\mathbb{E}\left[\beta \sum_{t=t'}^{\infty} \delta^t u_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right)\right],$$

Note, $\boldsymbol{\theta}^u$ contains the structural parameters defined in the flow utility function along with ψ_3 that are to be jointly estimated. Importantly, $\boldsymbol{\varepsilon}_{it}$ and $\boldsymbol{\theta}^u$ are only known to the agent. When the present-bias factor $\beta < 1$, this creates a wedge between the current self's evaluation of the present versus the future implying the consumer's preferences are time-inconsistent: the current self cannot commit their future self to follow a preferred plan. A sophisticated agent correctly anticipates this inconsistency and knows that their future selves will also apply the present-bias factor β when evaluating continuation values. The appropriate solution concept is therefore a Markov Perfect Equilibrium (MPE) of the intrapersonal game between the consumer's current and future selves.

X.3.1. *Assumptions*

I make five assumptions that are standard in the dynamic discrete choice literature which allow for a closed-form solution of the conditional choice probabilities and govern the form of the likelihood function.

Assumption AS (Additive Separability)**.** *The flow utility function is additively separable in the observable and unobservable components:*

$$u_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right) = \overline{u}_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right) + \varepsilon_{iljt}$$

where $\overline{u}_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right)$ *is the representative utility.*[59]

Assumption IID (i.i.d. Unobservables)**.** *The unobserved state vector* $\boldsymbol{\varepsilon}_{it} = \left(\varepsilon_{i1t}, \varepsilon_{i2t}, \dots, \varepsilon_{i,|\mathcal{J}|,t}\right)$ *are independent and identically distributed across individuals* i *and time periods* t *with*

[57] These details follow the framework presented in Aguirregabiria and Mira (2010) and Hortaçsu and Joo (2023) tailored to my specific problem, extended to accommodate sophisticated quasi-hyperbolic agents.

[58] This is because even though individuals do not literally plan forever, using an infinite-horizon formulation lets me naturally capture how today's consumption choices affect well-being throughout the remainder of life without imposing an abrupt cutoff at an arbitrary age; distant future effects are downweighted smoothly, avoiding artificial edge effects while still reflecting that current behavior has long-term consequences. Moreover, this remains consistent with the rational addiction literature first proposed in Becker and Murphy (1988).

[59] Additive separability allows for the idiosyncratic shocks to be conveniently integrated out when forming the conditional choice probabilities.

CDF $F_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{it})$.[60]

Assumption CI-Observed (Conditional Independence of Observed State Variables)**.** *Conditional on the current values of the decision* d_{ijt} *and the observable state variables* $\boldsymbol{x}_{it} = (h_{it}, a^f_{it}, a^s_{it}, a^{flav}_{it}, \boldsymbol{p}_{it})$*, next period observable state variables do not depend on current* $\boldsymbol{\varepsilon}_{it}$*, i.e.,*

$$F_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}; \boldsymbol{\theta}^f\right) = F_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right),$$

where $\boldsymbol{\theta}^f = (\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \Sigma)$ *is the parameter vector governing the shape of this distribution.*[61]

Assumption LOGIT (Type One Extreme Value Distribution)**.** *The unobserved state vector* $\boldsymbol{\varepsilon}_{it} = \left(\varepsilon_{i1t}, \varepsilon_{i2t}, \ldots, \varepsilon_{i,|\mathcal{J}|,t}\right)$ *are independent and identically distributed across alternatives* $j \in \mathcal{J}$ *with type one extreme value distribution.*[62]

Assumption DIS (Discrete Support of Observed State Variables)**.** *The support of* $\boldsymbol{x}_{it} = (a^f_{it}, a^s_{it}, a^{flav}_{it}, \boldsymbol{p}_{it})$ *is discrete and finite:* $\boldsymbol{x}_{it} \in \mathcal{X} \equiv \left\{\boldsymbol{x}^1_{it}, \boldsymbol{x}^2_{it}, \ldots, \boldsymbol{x}^{|\mathcal{X}|}_{it}\right\}$.[63]

Note that Assumptions CI-Observed and IID together imply that

$$\begin{aligned} F_{\boldsymbol{s}}\left(\boldsymbol{s}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{s}_{it}; \boldsymbol{\theta}^f\right) &= F_{\boldsymbol{x},\boldsymbol{\varepsilon}}\left(\boldsymbol{x}_{i,t+1}, \boldsymbol{\varepsilon}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}; \boldsymbol{\theta}^f\right) \\ &= F_{\boldsymbol{\varepsilon}}\left(\boldsymbol{\varepsilon}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{i,t+1}, \boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}; \boldsymbol{\theta}^f\right) F_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}; \boldsymbol{\theta}^f\right) \\ &= F_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{i,t+1}) F_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right), \end{aligned}$$

where the first equality follows from the Law of Total Probability. Thus, the state transitions are governed by the formula $F_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{i,t+1}) F_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it}; \boldsymbol{\theta}^f\right)$, where $\boldsymbol{\theta}^f$ is dropped from $F_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{i,t+1})$ as we do not care about the parameters governing this distribution. This represents the conditional independence assumption of Rust (1987). Moreover, these assumptions imply that the transition probabilities are both stationary and follow a first order Markov process.

X.3.2. *Sophisticated Agent Value Functions*

I define the decision utility, which captures how the agent evaluates alternatives when choosing (and therefore discounts the continuation value by $\beta\delta$):

$$\overline{v}_{dijt}\left(\boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) = \overline{u}_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right) + \beta\delta \int_{\mathcal{X}} V\left(\boldsymbol{x}_{i,t+1}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) \, dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right).$$

[60] While i.i.d. unobservables over individuals $i \in \{1, 2, \ldots, N\}$ is realistic in most cases, i.i.d. unobservables over time periods $t \in \{1, 2, \ldots\}$ is more concerning.

[61] This assumption implies that $\boldsymbol{\varepsilon}_{it}$ affects the next observed component of the state $\boldsymbol{x}_{i,t+1}$ only through the control variable $\boldsymbol{d}_{it}$.

[62] This assumption allows for the conditional choice probabilities to be represented in a closed-form.

[63] This assumption allows for Monte Carlo integration to be applied when solving for the value function.

I define the experienced utility, which captures the actual payoff that accrues from a given choice (and therefore discounts the continuation value by δ alone):

$$\overline{v}_{eijt}\left(\boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) = \overline{u}_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right) + \delta \int_{\mathcal{X}} V\left(\boldsymbol{x}_{i,t+1}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\, dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right).$$

The experienced utility $\overline{v}_{eijt}$ discounts by δ rather than $\beta\delta$ because β is a one-time wedge between "now" and "all future periods," not a per-period discount factor. Between any two consecutive future periods, the discount is simply δ. The distinction between $\overline{v}_{dijt}$ and $\overline{v}_{eijt}$ matters because the future self chooses according to $\overline{v}_{dijt}$ (which incorporates β) but experiences the payoff $\overline{v}_{eijt}$ (which does not).

Proposition 1 (Sophisticated Bellman representation)**.** *Let the future self choose*

$$j^* = \arg\max_{j \in \mathcal{J}} \left\{ \overline{v}_{dijt}\left(\boldsymbol{d}_{i,t+1}, \boldsymbol{x}_{i,t+1}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) + \varepsilon_{ij,t+1} \right\}.$$

Then the continuation value can be written as

$$V\left(\boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) = \sum_{j'=1}^{|\mathcal{J}|} p_{ilj't}\left(\boldsymbol{x}_{it}\right) \overline{v}_{eij't}\left(\boldsymbol{d}'_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) - \sum_{j'=1}^{|\mathcal{J}|} p_{ilj't}\left(\boldsymbol{x}_{it}\right) \log\left(p_{ilj't}\left(\boldsymbol{x}_{it}\right)\right),$$

where

$$p_{ilj't}\left(\boldsymbol{x}_{it}\right) = \frac{\exp\left[\overline{v}_{dij't}\left(\boldsymbol{d}'_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\right]}{\sum_{j''=1}^{|\mathcal{J}|} \exp\left[\overline{v}_{dij''t}\left(\boldsymbol{d}''_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\right]}.$$

Proof. The realized one-period payoff of the chosen alternative is

$$\overline{v}_{eij^*t} + \varepsilon_{ij^*,t+1}.$$

Taking expectation over $\boldsymbol{\varepsilon}_{i,t+1}$,

$$\mathbb{E}_{\boldsymbol{\varepsilon}_{i,t+1}}\left[\overline{v}_{eij^*t} + \varepsilon_{ij^*,t+1}\right] = \sum_{j'=1}^{|\mathcal{J}|} p_{ilj't}\left[\overline{v}_{eij't} + \mathbb{E}\left[\varepsilon_{ij',t+1} \mid j' \text{ chosen}\right]\right].$$

Under the Type I extreme value assumption,

$$\mathbb{E}\left[\varepsilon_{ij',t+1} \mid j' \text{ chosen}\right] = -\log p_{ilj't}$$

up to the Euler constant, which is dropped under the standard normalization.[64] Sub-

[64] Under T1EV, $\mathbb{E}[\varepsilon_j \mid j \text{ chosen}] = \gamma - \log p_j$. Dropping γ is standard in dynamic discrete choice because it is an additive constant.

stituting gives

$$\mathbb{E}_{\boldsymbol{\varepsilon}_{i,t+1}}\left[\overline{v}_{eij^*t}+\varepsilon_{ij^*,t+1}\right]=\sum_{j'=1}^{|\mathcal{J}|}p_{ilj't}\overline{v}_{eij't}-\sum_{j'=1}^{|\mathcal{J}|}p_{ilj't}\log p_{ilj't}.$$

This is the fixed-point mapping for V. □

Corollary 1 (Exponential discounting as a special case)**.** *When $\beta=1$, the decision utility and experienced utility coincide: $\overline{v}_{dijt}=\overline{v}_{eijt}$. In this case, the continuation value reduces to the standard logsumexp expression*

$$V\left(\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\boldsymbol{\theta}^f\right)=\log\left(\sum_{j'=1}^{|\mathcal{J}|}\exp\left\{\overline{v}_{eij't}\left(\boldsymbol{d}'_{it},\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\boldsymbol{\theta}^f\right)\right\}\right).$$

Proof. If $\beta=1$, then $\overline{v}_{dijt}=\overline{v}_{eijt}$ for all j. Hence $p_{ilj't}$ is the softmax of $\overline{v}_{eij't}$. Substituting this into the fixed-point expression

$$V=\sum_{j'=1}^{|\mathcal{J}|}p_{ilj't}\,\overline{v}_{eij't}-\sum_{j'=1}^{|\mathcal{J}|}p_{ilj't}\log p_{ilj't}$$

and using the softmax identity

$$\sum_{j=1}^{|\mathcal{J}|}p_j v_j-\sum_{j=1}^{|\mathcal{J}|}p_j\log p_j=\log\left(\sum_{j=1}^{|\mathcal{J}|}e^{v_j}\right)$$

yields the displayed logsumexp form. □

X.3.3. *Conditional Choice Probabilities*

The conditional choice probability, $\mathbb{P}_{ilj t}\left(\boldsymbol{d}_{it}\mid\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\boldsymbol{\theta}^f\right)$, represents the probability that individual i chooses alternative j at time period t, which is captured by $\boldsymbol{d}_{it}$, conditional on observable state information $\boldsymbol{x}_{it}$. It is obtained by integrating the optimal decision rule over the unobservable state variables:

$$\mathbb{P}_{ilj t}\left(\boldsymbol{d}_{it}\mid\boldsymbol{x}_{it};\boldsymbol{\theta}^u,\boldsymbol{\theta}^f\right)=\int_{\mathbb{R}^J}\mathbb{1}\left\{\overline{v}_{dijt}(\boldsymbol{d}_{it},\boldsymbol{x}_{it})+\varepsilon_{iljt}>\overline{v}_{dij't}\left(\boldsymbol{d}'_{it},\boldsymbol{x}_{it}\right)+\varepsilon_{ij't}\ \forall j'\neq j\right\}\ dF_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}_{it}).$$

This gives the probability mass of all shock vectors $\boldsymbol{\varepsilon}_{it}$ such that alternative j yields the highest decision utility, i.e., consuming product j at time period t is i's optimal action conditional on their decision utility $\overline{v}_{dijt}(\boldsymbol{d}_{it},\boldsymbol{x}_{it})$.[65] These conditional choice

[65] This is a J dimensional integral, with J representing the number of available alternatives exclusive alternative j. Importantly, this integral represents the probability that an individual i chooses alternative j at time t due the unobserved factors for j being sufficiently better than those for all other alternatives j' to overcome any advantage that j' has on observed factors. Thus, we are conditioning on alternative j and integrating over the remaining J alternatives.

probabilities have a closed-form expression of

$$\begin{aligned}\mathbb{P}_{iljt}\left(\boldsymbol{d}_{it} \mid \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) &= \frac{\exp\left[\overline{v}_{dijt}\left(\boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\right]}{\sum_{j' \in \mathcal{J}} \exp\left[\overline{v}_{dij't}\left(\boldsymbol{d}'_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\right]} \\ &= \frac{\exp\left\{\overline{u}_{iljt}\left(\boldsymbol{d}_{it}; \boldsymbol{\theta}^u\right) + \beta\delta \int_{\mathcal{X}} V\left(\boldsymbol{x}_{i,t+1}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\ dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right)\right\}}{\sum_{j' \in \mathcal{J}} \exp\left\{\overline{u}_{ij't}\left(\boldsymbol{d}'_{it}; \boldsymbol{\theta}^u\right) + \beta\delta \int_{\mathcal{X}} V\left(\boldsymbol{x}_{i,t+1}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right)\ dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}'_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right)\right\}},\end{aligned}$$

which follows from the assumptions placed on the random vector $\boldsymbol{\varepsilon}_{it}$ (McFadden (1972)).[66] Critically, the conditional choice probabilities are determined by the decision utility $\overline{v}_{dijt}$ (with $\beta\delta$ discounting), not the experienced utility $\overline{v}_{eijt}$. This is because the agent chooses the alternative that maximizes their present-biased evaluation $\overline{v}_{dijt} + \varepsilon_{iljt}$, even though the actual payoff from that choice is $\overline{v}_{eijt} + \varepsilon_{iljt}$. Importantly, notice that the continuation value $V\left(\boldsymbol{x}_{i,t+1}\right)$ appearing in both the numerator and denominator is the sophisticated continuation value derived from the fixed-point equation above, and the expectation over the observed state space captured by the integral over $\mathcal{X}$ can be evaluated via Monte Carlo integration, providing a mechanism to recover these conditional choice probabilities.

X.3.4. *Likelihood Function*

Since the household's latent type $l \in \{1, 2\}$ is unobserved, the full likelihood for household i integrates over both types. The conditional choice probabilities $\mathbb{P}_{iljt}$ are now determined by the type-specific decision utility $\overline{v}_{dljt}$ as derived above. The household-specific mixing weights are

$$\pi_{i2} = \sigma\left(\pi_0 + \pi_{\text{TYA}} \cdot \overline{h}_i\right),$$

where $\sigma(\cdot)$ is the sigmoid function, $\pi_{i1} = 1 - \pi_{i2}$, π_0 is a baseline intercept, π_{TYA} is the TYA share shifter, and $\overline{h}_i$ is the fraction of months household i has a teen or young adult present. The contribution of household i to the likelihood function is

$$\mathcal{L}_i\left(\boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) = \sum_{l=1}^{2} \pi_{il} \cdot \mathbb{P}\left(\boldsymbol{x}_{i1} \mid \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) \cdot \prod_{t=1}^{T_i} \mathbb{P}_{iljt}\left(\boldsymbol{d}_{it} \mid \boldsymbol{x}_{it}; \boldsymbol{\theta}^u, \boldsymbol{\theta}^f\right) \cdot \prod_{t=1}^{T_i - 1} f_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid \boldsymbol{d}_{it}, \boldsymbol{x}_{it}; \boldsymbol{\theta}^f\right),$$

where $\boldsymbol{d}_{it}$ is the choice vector of household i in period t. Inside the sum over types, the three components are: (i) the density of the initial state $\mathbb{P}(\boldsymbol{x}_{i1})$, (ii) the product of type-specific conditional choice probabilities across periods, and (iii) the product of state

[66] The integral used to obtain the conditional choice probabilities is calculated over $\mathbb{R}^J$ in general. However, after the assumption of i.i.d. type one extreme value shocks, the integral reduces to $\mathbb{R}$ and uses the continuous version of the Law of Total Probability to derive the closed-form expression.

transition densities. The same simplifications as in the original formulation apply: the initial-conditions density does not depend on $\boldsymbol{\theta}^u$ and drops, and the transition density term is fully determined by the first-stage estimates $\widehat{\boldsymbol{\theta}^f}$ and drops. What remains is a pseudo-likelihood based solely on the type-specific conditional choice probabilities:

$$\max_{\boldsymbol{\theta}^u} \mathscr{L}\left(\boldsymbol{\theta}^u, \widehat{\boldsymbol{\theta}^f}\right) = \max_{\boldsymbol{\theta}^u} \left\{ \sum_{i=1}^{N} \log \left[\sum_{l=1}^{2} \pi_{il} \cdot \prod_{t=1}^{T_i} \mathbb{P}_{iljt}\left(\boldsymbol{d}_{it} \mid \boldsymbol{x}_{it}\right) \right] \right\},$$

where $\mathscr{L}$ denotes the second-stage pseudo-log-likelihood. The inner product $\prod_{t=1}^{T_i} \mathbb{P}_{iljt}$ is the probability of observing household i's entire purchase history if it were type l. The inner sum integrates over the unobserved type by weighting each type's panel likelihood by the household-specific mixing probability π_{il}.

X.4. *Computational Details*

X.4.1. *Alternatives*

The choice set is discretized into $|\mathcal{J}| = 40$ alternatives. Alternative $j = 1$ is the outside option. The remaining alternatives are quantity-bin choices for cigarettes, original e-cigarettes, non-FDA authorized flavored e-cigarettes, FDA authorized flavored e-cigarettes, and bundles. Specifically, there are 12 cigarette alternatives, 7 original e-cigarette alternatives, 7 non-FDA authorized flavored e-cigarette alternatives, 7 FDA authorized flavored e-cigarette alternatives, and 6 bundle alternatives (2 per e-cigarette subcategory).

Each alternative corresponds to a bin, and the consumption assigned to that alternative is the empirical median quantity within that bin. For bins defined by a closed range, I use the median of observed quantities in that range. For singleton bins, the assigned consumption equals the singleton value itself. This mapping converts the high-dimensional product space into a finite set of representative quantity choices while preserving the center of each bin's empirical distribution.

Cigarette alternatives are defined over total packs with no e-cigarette purchase, using 12 bins: 1 pack, 2 packs, 3–4 packs, 5–9 packs, 10 packs, 11–19 packs, 20 packs, 21–29 packs, 30 packs, 31–39 packs, 40 packs, and 41+ packs. Each of the three e-cigarette subcategories (original, non-FDA authorized flavored, and FDA authorized flavored) uses the same 7 quantity bins defined over total e-liquid in milliliters with no cigarette purchase: 0–5 mL, 5–10 mL, 10–15 mL, 15–20 mL, 20–30 mL, 30–50 mL, and 50+ mL, where the left endpoint is exclusive and the right endpoint is inclusive.

Bundle alternatives are constructed by crossing e-cigarette subcategory (original, non-FDA authorized flavored, FDA authorized flavored) with a cigarette-intensity split (low: ≤ 20 packs versus high: > 20 packs), yielding six possible bundles. For each bundle alternative, cigarette and e-cigarette consumption components are assigned using

the empirical medians within that bundle cell, so each bundle carries two representative quantities, one for cigarettes and one for e-cigarettes.

X.4.2. *State Space*

At each time period t, consumer i observes the state vector $\boldsymbol{s}_{it} = (\boldsymbol{x}_{it}, \boldsymbol{\varepsilon}_{it})$, where $\boldsymbol{x}_{it} = \left(h_{it}, a^f_{it}, a^s_{it}, a^{\text{flav}}_{it}, \boldsymbol{p}_{it}\right)$ denotes the components of the state that are observable to the econometrician: the teen or young-adult indicator $h_{it} \in \{0, 1\}$, the fast and slow addiction stocks a^f_{it} and a^s_{it}, the flavored habit stock a^{flav}_{it}, and the vector of prices $\boldsymbol{p}_{it}$. To solve the dynamic program, the continuous components of $\boldsymbol{x}_{it}$ must be discretized to facilitate the evaluation of expectations over continuation values via Monte Carlo integration.

The teen or young-adult component is a binary indicator: $h_{it} = 1$ if at least one household member is aged 13–25 in period t, and $h_{it} = 0$ otherwise. Because h_{it} is a predetermined household characteristic that does not respond to nicotine purchasing decisions, it enters the flow utility as a covariate but does not require modeling transitions or integration over future demographic states in the continuation value.

As described in Section V.4., the model decomposes addiction into a fast stock a^f_{it} (craving) and a slow stock a^s_{it} (dependence), each evolving according to its own law of motion:

$$\begin{aligned} a^s_{i,t+1} &= (1 - \psi_1) \cdot a^s_{it} + n_{iCt} + n_{iEt}, \\ a^f_{i,t+1} &= (1 - \psi_2) \cdot a^f_{it} + n_{iCt} + n_{iEt}, \end{aligned}$$

where $\psi_1 = 0.10$ and $\psi_2 = 0.90$ are the decay rates for the slow and fast stocks, respectively, and nicotine inflow is bounded such that

$$0 \leq n_{iCt} + n_{iEt} \leq n_{\max}.$$

For each stock $m \in \{f, s\}$, assuming $0 < \psi_m \leq 1$, addiction depreciates geometrically over time. Iterating forward yields

$$a^m_{i,t+k} = (1 - \psi_m)^k a^m_{it} + \sum_{j=0}^{k-1} (1 - \psi_m)^{k-1-j} (n_{iC,t+j} + n_{iE,t+j}).$$

Using the upper bound on inflows implies

$$\begin{aligned} a^m_{i,t+k} &\leq (1 - \psi_m)^k a^m_{it} + n_{\max} \sum_{r=0}^{k-1} (1 - \psi_m)^r \\ &= (1 - \psi_m)^k a^m_{it} + \frac{n_{\max}}{\psi_m} \left(1 - (1 - \psi_m)^k\right). \end{aligned}$$

Taking limits as $k \to \infty$ and using $0 \leq 1 - \psi_m < 1$ gives

$$\limsup_{k\to\infty} a^m_{i,t+k} \leq \frac{n_{\max}}{\psi_m}.$$

Equivalently, if the consumer were to choose the maximal inflow $n_{iCt} + n_{iEt} = n_{\max}$ in every period, the steady-state addiction level a^{m*} satisfies

$$\begin{aligned} a^{m*} &= (1-\psi_m)a^{m*} + n_{\max}, \\ \psi_m a^{m*} &= n_{\max}, \\ a^{m*} &= \frac{n_{\max}}{\psi_m}. \end{aligned}$$

I define the normalized addiction stock used during estimation as $\tilde{a}^m_{it} \equiv \psi_m \, a^m_{it}$ and multiply both sides of the law of motion by ψ_m to obtain

$$\tilde{a}^m_{i,t+1} = (1-\psi_m)\,\tilde{a}^m_{it} + \psi_m\,(n_{iCt} + n_{iEt}).$$

Since $n_{iCt} + n_{iEt} \in [0, n_{\max}]$ and nicotine is standardized so that $n_{\max} = 1$, the steady-state upper bound is $\tilde{a}^{m*} = \psi_m \cdot (n_{\max}/\psi_m) = 1$, so $\tilde{a}^m_{it} \in [0,1]$ for all t. Consequently, I discretize each normalized addiction state space as equally spaced grid points in the unit interval:

$$\begin{aligned} \mathcal{A}^f &= \left\{0 = \tilde{a}^f_1, \ldots, \tilde{a}^f_5 = 1\right\}, \\ \mathcal{A}^s &= \left\{0 = \tilde{a}^s_1, \ldots, \tilde{a}^s_{10} = 1\right\}, \\ \mathcal{A}^{\text{flav}} &= \left\{0 = \tilde{a}^{\text{flav}}_1, \ldots, \tilde{a}^{\text{flav}}_{10} = 1\right\}, \end{aligned}$$

with $|\mathcal{A}^f| = 5$ grid points for the fast stock, $|\mathcal{A}^s| = 10$ grid points for the slow stock, and $|\mathcal{A}^{\text{flav}}| = 10$ grid points for the flavored habit stock. The slow stock receives a finer grid than the fast stock because its low decay rate ($\psi_1 = 0.10$) makes it the primary driver of long-run addiction dynamics. For notational brevity, I drop the tilde hereafter; all subsequent references to the addiction states (including the value function iteration grid and the laws of motion in the VFI) refer to the normalized stocks $\tilde{a}^f$, $\tilde{a}^s$, and $\tilde{a}^{\text{flav}}$.

Focusing now on discretizing the pricing space, the dynamic state tracks category-level per-unit prices for cigarettes and e-cigarettes. Specifically, I construct per-unit prices from the data and then form the empirical distributions for cigarettes and e-cigarettes separately. The price grid uses seven quantile points per category (5th, 20th, 35th, 50th, 65th, 80th, and 95th percentiles), giving

$$\begin{aligned} \mathcal{P}_{\text{cig}} &= \left\{p^{\text{cig}}_1, \ldots, p^{\text{cig}}_7\right\}, \\ \mathcal{P}_{\text{ecig}} &= \left\{p^{\text{ecig}}_1, \ldots, p^{\text{ecig}}_7\right\}. \end{aligned}$$

The combined price-state grid used in the dynamic program is the Cartesian product $\mathcal{P}$ of these two one-dimensional grids, yielding $|\mathcal{P}| = |\mathcal{P}_{\text{cig}}| \times |\mathcal{P}_{\text{ecig}}| = 49$ price states.

It is important to account for non-linear pricing schedules, i.e., consumers often receive discounts for buying in larger quantities. I investigate these non-linearities by plotting the distribution of per-unit prices paid across different purchase quantities for each category in Figure A4.

Figure A4. Per-Pack of Cigarettes (Panel (a)) and Per-Milliliter of E-Liquid (Panel (b)) Prices Paid.

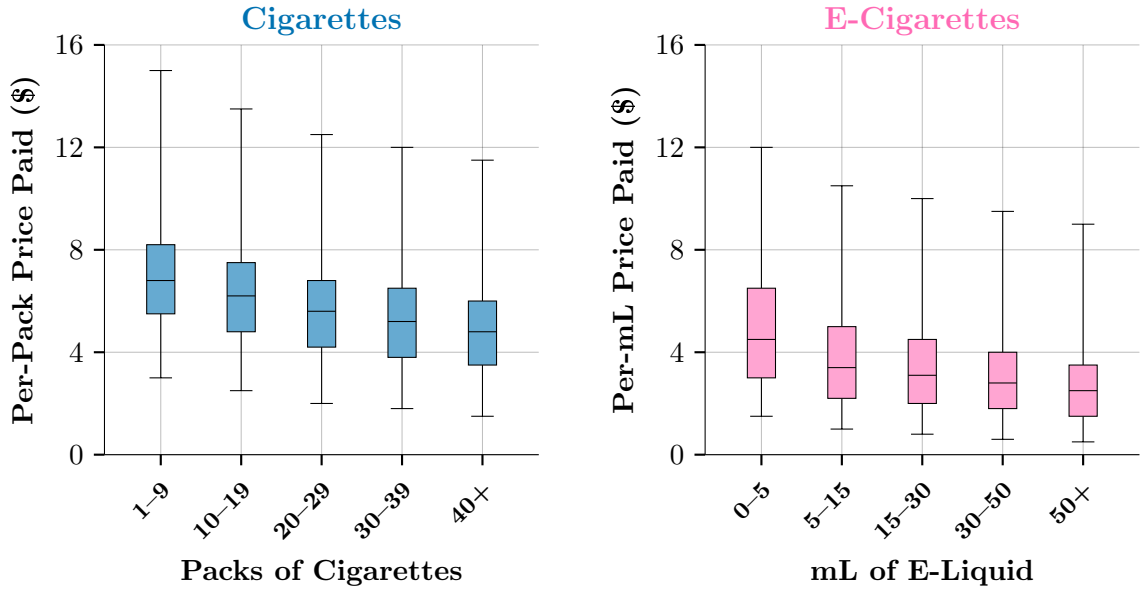

The black horizontal line represents the median, the shaded boxes represent the interquartile range, and the lower and upper gray horizontal lines represent the minimum and maximum, respectively.

Nonlinear pricing schedules are handled outside the state discretization through alternative-specific price ratios. The price state stores a category-level benchmark per-unit price, while a ratio term adjusts that benchmark up or down based on the quantity profile of alternative j. The expenditure term entering utility is

$$E[p, j] = p_{\text{cig}}[p] \cdot \text{ratio}_{\text{cig}}[j] \cdot q_{\text{cig}}[j] + p_{\text{ecig}}[p] \cdot \text{ratio}_{\text{ecig}}[j] \cdot q_{\text{ecig}}[j].$$

Each ratio is computed as the median per-unit price within a quantity bin divided by the overall median per-unit price for that category. Small-quantity alternatives typically have ratios above one (higher per-unit cost) and large-quantity alternatives below one (quantity discounts). For bundle alternatives, each component is mapped separately to the nearest standalone bin; for components not consumed in alternative j, the ratio enters as one.[67]

[67] For example, suppose the cigarette benchmark price is $p_{\text{cig}}[p] = 8$. A small-quantity alternative with $\text{ratio}_{\text{cig}}[j] = 1.10$ and $q_{\text{cig}}[j] = 2$ yields expenditure $8 \cdot 1.10 \cdot 2 = 17.6$ (per-unit cost of 8.8), while

Since the teen or young-adult indicator takes $|\mathcal{H}| = 2$ values, the fast addiction grid contains $|\mathcal{A}^f| = 5$ points, the slow addiction grid contains $|\mathcal{A}^s| = 10$ points, the flavored habit grid contains $|\mathcal{A}^{\text{flav}}| = 10$ points, and the pricing state space contains $|\mathcal{P}| = 49$ points, the full observable state space $\mathcal{X}$ consists of

$$|\mathcal{X}| = |\mathcal{H}| \times |\mathcal{A}^f| \times |\mathcal{A}^s| \times |\mathcal{A}^{\text{flav}}| \times |\mathcal{P}| = 2 \times 5 \times 10 \times 10 \times 49 = 49{,}000$$

distinct states.

X.4.3. *Data Standardization*

Consumption quantities and nicotine entering the flow utility are standardized by their sample maxima before estimation. This ensures that the optimizer works with parameters of comparable scale, improving numerical stability. Prices are not standardized; they enter in real dollar units directly. Let $\overline{q}_C$, $\overline{q}_E$, $\overline{q}_{CE}$, and $\overline{n}$ denote the sample maxima of cigarette quantity (packs), e-cigarette quantity (mL), bundle interaction quantity (packs $\times$ mL), and nicotine content (mg), respectively. The standardized variables entering the flow utility are

$$\begin{aligned}
\widetilde{q}_{\text{cig}}[j] &= q_{\text{cig}}[j] \,/\, \overline{q}_C, \\
\widetilde{q}_{\text{ecig}}[j] &= q_{\text{ecig}}[j] \,/\, \overline{q}_E, \\
\widetilde{q}_{\text{bundle}}[j] &= q_{\text{cig}}[j] \cdot q_{\text{ecig}}[j] \,/\, \overline{q}_{CE}, \\
\widetilde{n}[j] &= n[j] \,/\, \overline{n}.
\end{aligned}$$

The nicotine standardization also affects the addiction stocks. Because the normalized addiction stock evolves as $\widetilde{a}^m_{t+1} = (1 - \psi_m)\widetilde{a}^m_t + \psi_m \widetilde{n}[j_t]$ for $m \in \{f, s\}$, both stocks converge to $\widetilde{n}[j]$ at steady state regardless of ψ_m. Consequently, the composite stock $a = (\widetilde{a}^f + \widetilde{a}^s)/2$ equals $\widetilde{n}[j] = n[j]/\overline{n}$ at steady state, so the addiction level is implicitly scaled by $1/\overline{n}$ relative to raw nicotine units.

The estimated parameters are in standardized units. To recover the original-unit parameters, the following rescaling is applied:

$$\begin{aligned}
\widehat{\alpha_C}^{\text{orig}} &= \widehat{\alpha}_C \,/\, \overline{q}_C, \\
\widehat{\alpha_E}^{\text{orig}} &= \widehat{\alpha}_E \,/\, \overline{q}_E, \\
\widehat{\alpha_{CE}}^{\text{orig}} &= \widehat{\alpha}_{CE} \,/\, \overline{q}_{CE}, \\
\widehat{\omega_C}^{\text{orig}} &= \widehat{\omega}_C \,/\, \overline{q}_C, \\
\widehat{\omega_E}^{\text{orig}} &= \widehat{\omega}_E \,/\, \overline{q}_E, \\
\widehat{\gamma_1}^{\text{orig}} &= \widehat{\gamma}_1 \,/\, \overline{n}.
\end{aligned}$$

a large-quantity alternative with $\text{ratio}_{\text{cig}}[j'] = 0.90$ and $q_{\text{cig}}[j'] = 20$ yields $8 \cdot 0.90 \cdot 20 = 144$ (per-unit cost of 7.2).

Since prices enter in original units (real dollars per unit), ω_C and ω_E require only rescaling by the quantity maxima.

The flavored habit stock requires a separate derivation. In the paper, the law of motion is written in the accumulating form

$$s_{t+1}^{\text{flav}} = (1 - \psi_3)\, s_t^{\text{flav}} + \mathbb{1}\{j_t \in \mathcal{E}_f\},$$

where the steady state under perpetual flavored purchasing is $s_\infty^{\text{flav}} = 1/\psi_3$. To keep the state variable on $[0, 1]$ for numerical stability, estimation uses the normalized form $\widetilde{a}^{\text{flav}} = \psi_3 \cdot s^{\text{flav}}$, which evolves as

$$\widetilde{a}_{t+1}^{\text{flav}} = (1 - \psi_3)\, \widetilde{a}_t^{\text{flav}} + \psi_3 \cdot \mathbb{1}\{j_t \in \mathcal{E}_f\}.$$

To verify: multiply both sides of the paper's law of motion by ψ_3 and substitute $\widetilde{a}^{\text{flav}} = \psi_3 \cdot s^{\text{flav}}$:

$$\begin{aligned}
\psi_3\, s_{t+1}^{\text{flav}} &= (1 - \psi_3)\, \psi_3\, s_t^{\text{flav}} + \psi_3 \cdot \mathbb{1}\{j_t \in \mathcal{E}_f\} \\
\widetilde{a}_{t+1}^{\text{flav}} &= (1 - \psi_3)\, \widetilde{a}_t^{\text{flav}} + \psi_3 \cdot \mathbb{1}\{j_t \in \mathcal{E}_f\}.
\end{aligned}$$

The estimated $\widehat{\gamma}_2$ and $\widehat{\gamma}_3$ multiply $\widetilde{a}^{\text{flav}}$ in the flow utility. Since $\widetilde{a}^{\text{flav}} = \psi_3 \cdot s^{\text{flav}}$, equating the utility contributions gives

$$\widehat{\gamma}_2 \cdot \widetilde{a}_t^{\text{flav}} = \widehat{\gamma}_2 \cdot \psi_3 \cdot s_t^{\text{flav}} = \gamma_2^{\text{orig}} \cdot s_t^{\text{flav}},$$

from which

$$\begin{aligned}
\gamma_2^{\text{orig}} &= \widehat{\gamma}_2 \times \psi_3, \\
\gamma_3^{\text{orig}} &= \widehat{\gamma}_3 \times \psi_3, \\
\gamma_4^{\text{orig}} &= \widehat{\gamma}_4 \times \psi_3.
\end{aligned}$$

The remaining parameters require no rescaling: λ_1–λ_4 multiply binary indicators, ξ_{lk} are additive fixed effects, and the mixing weight parameters π_0 and π_{TYA} enter the logistic mixing function directly.

X.4.4. *AR(1) Price Process Estimation*

The evolution of category-level per-unit prices is modeled as an autoregressive process of order one (AR(1)) taking the form

$$\begin{aligned}
p_{C,t+1} &= \phi_{0C} + \phi_{1C} \cdot p_{Ct} + \eta_{C,t+1}, \\
p_{E,t+1} &= \phi_{0E} + \phi_{1E} \cdot p_{Et} + \eta_{E,t+1},
\end{aligned}$$

with

$$\begin{bmatrix} \eta_{C,t+1} \\ \eta_{E,t+1} \end{bmatrix} \sim \mathbb{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Sigma \right),$$

where p_{Ct} and p_{Et} denote the per-unit prices for cigarettes and e-cigarettes at time t, ϕ_{0C} and ϕ_{0E} are category-specific intercepts, ϕ_{1C} and ϕ_{1E} capture price persistence, and $\eta_{C,t+1}$ and $\eta_{E,t+1}$ are mean-zero normally distributed price shocks with covariance matrix Σ. This process operates at the category level: all alternatives $j \in \mathcal{J}_k$ share the same price parameters.[68] Notably, consumers form expectations over future prices at the category level rather than the product level, substantially reducing the dimensionality of the state space.[69] The AR(1) is estimated separately for each category, so price expectations are category-specific.

To implement this, I first compute product-level per-unit prices and then aggregate to monthly category-level medians. The AR(1) is estimated separately for each category on these median per-unit price series to obtain $\left(\widehat{\phi}_{0C}, \widehat{\phi}_{1C}\right)$ and $\left(\widehat{\phi}_{0E}, \widehat{\phi}_{1E}\right)$. The residual covariance matrix $\widehat{\Sigma}$ is then estimated from the cross-category residuals. These estimates allow me to simulate price transitions which are used to compute expected continuation values via Monte Carlo integration in the dynamic model.

X.4.5. *Transition Probabilities*

The only transition object constructed in the first stage is for prices. The teen-or-young-adult indicator $h_{it} \in \{0, 1\}$ is treated as a fixed household covariate and does not require a transition process.

For prices, I estimate category-level AR(1) processes for cigarettes and e-cigarettes and then simulate next-period prices with quasi-Monte Carlo integration. Let $m \in \{1, \ldots, |\mathcal{P}|\}$ index current combined price states and $r \in \mathcal{R} = \{1, \ldots, |\mathcal{R}|\}$ index Halton draws, with $|\mathcal{R}| = 200$. I generate correlated shocks by first drawing Halton normals $\boldsymbol{z}_r \sim \mathbb{N}(\mathbf{0}, I_2)$. I use the variance-covariance matrix $\widehat{\Sigma}$ obtained by the AR(1) process to recover the estimated residual vector for each Halton draw r as

$$\widehat{\boldsymbol{\eta}}_r = \widehat{L}\boldsymbol{z}_r,$$

where $\widehat{L}\widehat{L}' = \widehat{\Sigma}$. This implies $\widehat{\boldsymbol{\eta}}_r \sim \mathbb{N}\left(\mathbf{0}, \widehat{\Sigma}\right)$.

For each price state m and Halton draw r, I compute the predicted next-period

[68] Σ allows for correlation in price shocks across categories.

[69] This reduces the state space from $|\mathcal{A}^f| \times |\mathcal{A}^s| \times |\mathcal{P}|^J$ to $|\mathcal{A}^f| \times |\mathcal{A}^s| \times |\mathcal{P}|^K$, where the number of categories K is considerably smaller than the number of alternatives J.

prices as

$$\widehat{p}'_{C,m,r} = \widehat{\phi}_{0C} + \widehat{\phi}_{1C} \cdot p_{C,m} + \widehat{\eta}_{C,r},$$
$$\widehat{p}'_{E,m,r} = \widehat{\phi}_{0E} + \widehat{\phi}_{1E} \cdot p_{E,m} + \widehat{\eta}_{E,r}.$$

This yields a transition array with dimensions $|\mathcal{P}| \times |\mathcal{R}| \times 2$, where the last dimension corresponds to cigarettes and e-cigarettes.

The teen or young-adult indicator $h_{it} \in \{0, 1\}$ is a predetermined household characteristic that enters the flow utility but does not require transition modeling. Because h_{it} does not affect the continuation value, no integration over future demographic states is needed in the Bellman equation. The value function is solved separately for each value of h_{it}, conditioning on the household's observed TYA status.

X.4.6. *Initial Conditions for Addiction*

A central challenge in estimating dynamic models with persistent state variables is the treatment of initial conditions. In the present setting, both addiction stocks $a^f_{i,t}$ and $a^s_{i,t}$ are unobserved at the time a household first enters the panel. Because addiction enters utility directly and evolves dynamically over time, the specification of initial addiction levels may affect the likelihood.

Two features of the model substantially simplify this problem. First, each addiction stock evolves deterministically according to a linear law of motion. Second, nicotine consumption is itself a deterministic function of observed purchase choices. Together, these features imply that conditional on initial addiction levels $a^f_{i,0}$ and $a^s_{i,0}$ and observed consumption $\{n_{iC,t} + n_{iE,t}\}_{t=1}^{T_i}$, the entire addiction trajectory for household i is uniquely determined. Because the two stocks share the same nicotine inflow but differ only in their decay rates, the fixed-point procedure described below applies independently to each stock $m \in \{f, s\}$ with its own decay parameter.

Rather than imposing an ad hoc assumption on the initial conditions, I recover household-specific initial addiction stocks using a fixed-point procedure. The approach selects the initial addiction level for each household so that the addiction trajectory implied by observed behavior is internally consistent with the assumed initial condition.

Proposition 2 (Existence and uniqueness of initial addiction stock)**.** *Fix a household i with observed nicotine sequence $\{n_{iC,t} + n_{iE,t}\}_{t=1}^{T_i}$ and addiction depreciation rate $\psi_m \in (0, 1)$. Define the mapping*

$$F(a_{i,0}) \equiv a_{i,T_i}(a_{i,0}),$$

where $a_{i,T_i}(a_{i,0})$ denotes the terminal addiction stock obtained by simulating the addiction law of motion forward from initial condition $a_{i,0}$ using observed nicotine consumption.

Then:

1. *The mapping F admits a unique fixed point $a_{i,0}^{fp}$ satisfying*

$$a_{i,0}^{fp} = F\left(a_{i,0}^{fp}\right).$$

2. *For any initial guess $a_{i,0}^{(0)} \in \mathbb{R}$, the sequence defined by*

$$a_{i,0}^{(k+1)} = F\left(a_{i,0}^{(k)}\right)$$

converges globally and geometrically to $a_{i,0}^{fp}$.

Proof of Proposition 3. Fix a household i with observed nicotine sequence $\{n_{iC,t} + n_{iE,t}\}_{t=1}^{T_i}$, where T_i denotes the final observed month for household i. The normalized addiction stock evolves according to

$$a_{i,t} = (1-\psi)a_{i,t-1} + \psi\,(n_{iC,t} + n_{iE,t}),$$

where $t \in \{1, \ldots, T_i\}$. Iterating forward yields the closed-form expression

$$a_{i,T_i}(a_{i,0}) = (1-\psi)^{T_i} a_{i,0} + \psi \sum_{s=1}^{T_i} (1-\psi)^{T_i - s}(n_{iC,s} + n_{iE,s}).$$

Define the update operator $F : \mathbb{R} \to \mathbb{R}$ by

$$F(a_{i,0}) \equiv a_{i,T_i}(a_{i,0}) = (1-\psi)^{T_i} a_{i,0} + \psi \sum_{s=1}^{T_i} (1-\psi)^{T_i - s}(n_{iC,s} + n_{iE,s}).$$

For any $x, y \in \mathbb{R}$,

$$\begin{aligned} |F(x) - F(y)| &= \left|(1-\psi)^{T_i}(x-y)\right| \\ &= (1-\psi)^{T_i}|x-y|. \end{aligned}$$

As $\psi \in (0,1)$, it follows that $0 < (1-\psi)^{T_i} < 1$. So, F is a contraction mapping with Lipschitz constant

$$L_i \equiv (1-\psi)^{T_i}.$$

Since $\mathbb{R}$ is a complete metric space and F is a contraction, Banach's Fixed Point Theorem implies the existence and uniqueness of a fixed point $a_{i,0}^{fp}$ satisfying $a_{i,0}^{fp} = F\left(a_{i,0}^{fp}\right)$. Moreover, for any initial guess $a_{i,0}^{(0)} \in \mathbb{R}$, the fixed-point iteration $a_{i,0}^{(k+1)} = F\left(a_{i,0}^{(k)}\right)$ converges geometrically as

$$\left|a_{i,0}^{(k)} - a_{i,0}^{fp}\right| \leq L_i^k \left|a_{i,0}^{(0)} - a_{i,0}^{fp}\right|.$$

Solving $a_{i,0}^{fp} = F\left(a_{i,0}^{fp}\right)$ yields

$$\begin{aligned}
a_{i,0}^{fp} = F\left(a_{i,0}^{fp}\right) &\iff a_{i,0}^{fp} = (1-\psi)^{T_i} a_{i,0}^{fp} + \psi \sum_{s=1}^{T_i} (1-\psi)^{T_i - s}(n_{iC,s} + n_{iE,s}) \\
&\iff a_{i,0}^{fp} - (1-\psi)^{T_i} a_{i,0}^{fp} = \psi \sum_{s=1}^{T_i} (1-\psi)^{T_i - s}(n_{iC,s} + n_{iE,s}) \\
&\iff \left(1 - (1-\psi)^{T_i}\right) a_{i,0}^{fp} = \psi \sum_{s=1}^{T_i} (1-\psi)^{T_i - s}(n_{iC,s} + n_{iE,s}) \\
&\iff a_{i,0}^{fp} = \frac{\psi \sum_{s=1}^{T_i} (1-\psi)^{T_i - s}(n_{iC,s} + n_{iE,s})}{1 - (1-\psi)^{T_i}}.
\end{aligned}$$

□

The rate of convergence of this fixed-point procedure is governed by the addiction decay parameter ψ_m for $m \in \{s, f\}$. The influence of the initial condition on the period t addiction stock decays geometrically at rate $(1-\psi_m)^t$. After T_i observed periods, the initial condition's contribution to the terminal addiction level is $(1-\psi_m)^{T_i} a_{i,0}^m$, which vanishes rapidly for large ψ_m. The fast stock ($\psi_2 = 0.90$) converges almost immediately: $(1-0.90)^{T_i} = 0.10^{T_i}$, which falls below 10^{-3} within three months. The slow stock ($\psi_1 = 0.10$) converges more gradually: $(1-0.10)^{T_i} = 0.90^{T_i}$, requiring roughly 66 months for the factor to fall below 10^{-3}, but the fixed-point iteration still converges geometrically and yields a unique solution for any panel length.

X.4.7. *Initial Parameter Values*

The parameter estimates are updated during each iteration of the evaluation of the likelihood function. This implies an initial guess of the parameters $\boldsymbol{\theta}^u$ is required. After estimation, I then reinitialize the starting values to those achieved upon convergence and re-estimate the model once more, which results in the estimates obtained in Table 4.

X.4.8. *Flow Utilities*

I precompute flow utility on the full discrete state grid before each VFI solve for the current iteration of $\boldsymbol{\theta}^u$. The resulting array for type l is

$$\overline{u}_l \left[h, j, a^f, a^s, a^{\text{flav}}, p\right],$$

with dimensions $|\mathcal{H}| \times |\mathcal{J}| \times |\mathcal{A}^f| \times |\mathcal{A}^s| \times |\mathcal{A}^{\text{flav}}| \times |\mathcal{P}|$. This avoids recomputing static utility components inside the inner VFI loops.

For each state-alternative tuple, I evaluate the deterministic utility component

$$\begin{aligned}\overline{u}_l[h, j, a^f, a^s, a^{\text{flav}}, p] = \alpha_C \cdot \widetilde{q}_{\text{cig}}[j] + \alpha_E \cdot \widetilde{q}_{\text{ecig}}[j] + \alpha_{CE} \cdot \widetilde{q}_{\text{bundle}}[j] \\ + (\lambda_1 + \lambda_2 \cdot \mathbb{1}\{h = 1\}) \cdot \mathbb{1}\{j \in \mathcal{E}_f\} \\ + (\lambda_3 + \lambda_4 \cdot \mathbb{1}\{h = 1\}) \cdot \mathbb{1}\{j \in \mathcal{E}_{f,a}\} \\ + \gamma_1 \cdot a \cdot \mathbb{1}\{j = \mathcal{O}\} \\ + \gamma_2 \cdot a^{\text{flav}} \cdot \mathbb{1}\{j \in \mathcal{E}_o\} \\ + \gamma_3 \cdot a^{\text{flav}} \cdot \mathbb{1}\{j \in \mathcal{C}\} \\ + \gamma_4 \cdot a^{\text{flav}} \cdot \mathbb{1}\{j = \mathcal{O}\} \\ + \omega_C \cdot p_C[p] \cdot \widetilde{q}_{\text{cig}}[j] + \omega_E \cdot p_E[p] \cdot \widetilde{q}_{\text{ecig}}[j] \\ + \xi_{lk(j)},\end{aligned}$$

where $a = (a^f + a^s)/2$ is the unweighted average of the fast and slow addiction stocks, $\mathcal{E}_{f,a}$ denotes FDA-authorized flavored alternatives, $\mathcal{E}_o$ denotes the original (nonflavored) e-cigarette and original bundle alternatives, and $\mathcal{C}$ denotes cigarettes and original bundles, comprising products containing cigarettes without a flavored e-cigarette component. The γ_2 and γ_3 terms capture flavored habit lock-in: a high flavored habit stock penalizes switching to nonflavored e-cigarette products or cigarette-containing products, respectively. The γ_4 term captures flavored withdrawal: a high flavored habit stock penalizes market exit.

X.4.9. *Solving For the Value Function*

This subsection corresponds directly to the implemented routine. The algorithm is Jacobi VFI over states $(h, a^f, a^s, a^{\text{flav}}, p)$ with $h \in \mathcal{H}$, $a^f \in \mathcal{A}^f$, $a^s \in \mathcal{A}^s$, $a^{\text{flav}} \in \mathcal{A}^{\text{flav}}$, and $p \in \mathcal{P}$. The objects solved for are the value function $V \in \mathbb{R}^{|\mathcal{H}|\times|\mathcal{A}^f|\times|\mathcal{A}^s|\times|\mathcal{A}^{\text{flav}}|\times|\mathcal{P}|}$, the choice-specific decision value function $V_d \in \mathbb{R}^{|\mathcal{H}|\times|\mathcal{J}|\times|\mathcal{A}^f|\times|\mathcal{A}^s|\times|\mathcal{A}^{\text{flav}}|\times|\mathcal{P}|}$, and the experience choice-specific value function $V_e \in \mathbb{R}^{|\mathcal{H}|\times|\mathcal{J}|\times|\mathcal{A}^f|\times|\mathcal{A}^s|\times|\mathcal{A}^{\text{flav}}|\times|\mathcal{P}|}$. Because each latent type $l \in \{1, 2\}$ has its own category fixed effects $\xi_{lk(j)}$, these objects are computed separately for each type, and both types are solved in parallel.

Because I fix $\beta = 1$ during estimation, the decision and experienced utilities coincide ($V_{dl} = V_{el}$) and the distinction between them is immaterial for the estimated model. The separation only becomes important in the counterfactual simulations, where I vary $\beta < 1$ to explore the welfare implications of present bias.

I now state the exact update steps for a given type l:

1. Initialize $V_{l,\text{now}}$, $V_{l,\text{next}}$, $\overline{v}_{dl}$, and $\overline{v}_{el}$ to zeros.

2. For each state $(h, a^f, a^s, a^{\text{flav}}, p) \in \mathcal{H} \times \mathcal{A}^f \times \mathcal{A}^s \times \mathcal{A}^{\text{flav}} \times \mathcal{P}$ and alternative $j \in \mathcal{J}$:

(a) Compute next-period addiction stocks implied by choosing j:

$$\begin{aligned} a_j^{s\prime} &= (1-\psi_1)\cdot a^s + \psi_1 \cdot \widetilde{n}_j, \\ a_j^{f\prime} &= (1-\psi_2)\cdot a^f + \psi_2 \cdot \widetilde{n}_j, \\ a_j^{\text{flav}\prime} &= (1-\psi_3)\cdot a^{\text{flav}} + \psi_3 \cdot \mathbb{1}\left\{j \in \mathcal{E}_f\right\}, \end{aligned}$$

where $\widetilde{n}_j$ is the standardized nicotine content of alternative j. Because $a_j^{f\prime}$, $a_j^{s\prime}$, and $a_j^{\text{flav}\prime}$ generally fall between grid points, all continuation value lookups below use linear interpolation along each addiction and habit dimension. The interpolation brackets and weights for every reachable next-period stock are precomputed and stored prior to VFI to save during each VFI solve.

(b) For each Halton draw $r \in \{1, \ldots, |\mathcal{R}|\}$, compute predicted next-period prices $p'_{C,r}$ and $p'_{E,r}$ from the AR(1) process. These also generally fall between grid points, so bilinear interpolation over the price grid is used. As with addiction, the price brackets and weights for every draw are precomputed prior to VFI to save time during each VFI solve.

(c) Evaluate $V_{l,\text{now}}$ at the implied next-period state $(h, a_j^{f\prime}, a_j^{s\prime}, a_j^{\text{flav}\prime}, p'_{C,r}, p'_{E,r})$ for each draw r, using the interpolation objects from (a) and (b). Because the TYA indicator h is fixed within household (it does not transition), no integration over demographic states is needed. Average over draws to form the price-integrated continuation value:

$$\mathbb{E}\left[V_{l,\text{now}} \mid h, j, a^f, a^s, a^{\text{flav}}, p\right] = \frac{1}{|\mathcal{R}|}\sum_{r=1}^{|\mathcal{R}|} V_{l,\text{now}}\left(h, a_j^{f\prime}, a_j^{s\prime}, a_j^{\text{flav}\prime}, p'_{C,r}, p'_{E,r}\right).$$

(d) Compute choice-specific decision and experienced values:

$$\begin{aligned} \overline{v}_{dl}[h, j, a^f, a^s, a^{\text{flav}}, p] &= \overline{u}_l[h, j, a^f, a^s, a^{\text{flav}}, p] + \beta\delta \cdot \mathbb{E}\left[V_{l,\text{now}} \mid h, j, a^f, a^s, a^{\text{flav}}, p\right], \\ \overline{v}_{el}[h, j, a^f, a^s, a^{\text{flav}}, p] &= \overline{u}_l[h, j, a^f, a^s, a^{\text{flav}}, p] + \delta \cdot \mathbb{E}\left[V_{l,\text{now}} \mid h, j, a^f, a^s, a^{\text{flav}}, p\right]. \end{aligned}$$

3. For each $(h, a^f, a^s, a^{\text{flav}}, p)$, compute choice probabilities from the decision values:

$$\mathbb{P}_{iljt}[h, a^f, a^s, a^{\text{flav}}, p] = \frac{\exp\left(\overline{v}_{dl}[h, j, a^f, a^s, a^{\text{flav}}, p]\right)}{\sum_{j' \in \mathcal{J}} \exp\left(\overline{v}_{dl}[h, j', a^f, a^s, a^{\text{flav}}, p]\right)}.$$

4. Update the ex-ante value function:

$$V_{l,\text{next}}[h, a^f, a^s, a^{\text{flav}}, p] = \sum_{j \in \mathcal{J}} \mathbb{P}_{iljt}[h, a^f, a^s, a^{\text{flav}}, p] \cdot \overline{v}_{el}[h, j, a^f, a^s, a^{\text{flav}}, p] - \sum_{j \in \mathcal{J}} \mathbb{P}_{iljt}[h, a^f, a^s, a^{\text{flav}}, p]$$

This is the key sophisticated agent step: probabilities come from $\overline{v}_{dl}$ while payoffs

come from $\overline{v}_{el}$.

5. Convergence check:

$$\Delta_l = \max_{h,a^f,a^s,a^{\text{flav}},p} \left| V_{l,\text{next}}[h, a^f, a^s, a^{\text{flav}}, p] - V_{l,\text{now}}[h, a^f, a^s, a^{\text{flav}}, p] \right|.$$

   Set $V_{l,\text{now}} \leftarrow V_{l,\text{next}}$ and return to step 2 unless $\Delta_l < \tau$.

6. Post-convergence recomputation: recompute $\overline{v}_{dl}$ one more time from converged $V_{l,\text{now}}$ using the same Bellman decision-utility equation.

I use Jacobi iteration: all states are updated from the previous iterate $V_{l,\text{now}}$ before the full array is replaced. This makes each state update conditionally independent within an iteration, which is what enables parallelizing over $(a^s, a^{\text{flav}}, p)$.

X.4.10. *Likelihood Evaluation*

For each candidate $\boldsymbol{\theta}^u$, the VFI routine returns the type-specific decision-utility arrays

$$V_{dl}[h, j, a^f, a^s, a^{\text{flav}}, p],$$

for $l \in \{1, 2\}$, where $h \in \mathcal{H}$, $j \in \mathcal{J}$, $a^f \in \mathcal{A}^f$, $a^s \in \mathcal{A}^s$, $a^{\text{flav}} \in \mathcal{A}^{\text{flav}}$, and $p \in \mathcal{P}$. Because each latent type has its own category fixed effects $\xi_{lk(j)}$, the value function iteration is performed separately for each type, and both types are solved in parallel.

For each observation i, I observe $\left(y_i,\ h_i,\ a_i^f,\ a_i^s,\ a_i^{\text{flav}},\ p_{C,i},\ p_{E,i}\right)$. Here, y_i is the chosen alternative, h_i is the observed teen or young-adult state index, and $(a_i^f, a_i^s, a_i^{\text{flav}}, p_{C,i}, p_{E,i})$ are the continuous states to be interpolated over. The two continuous prices are representative category-level prices.

The interpolation step is performed for each type l:

1. Fix the demographic index at the observed value h_i. This means no interpolation is done over h.

2. Interpolate only over all addiction and habit stocks and prices inside the slice

$$V_{dl}[h_i, \cdot, \cdot, \cdot, \cdot, \cdot].$$

3. Use linear interpolation in each addiction and habit stock and bilinear interpolation in prices (cig, ecig), giving five-dimensional interpolation overall $(a^f, a^s, a^{\text{flav}}, p_C, p_E)$.

4. Obtain interpolated decision utilities for all alternatives for observation i and type l:

$$\widetilde{V}_l[i][j]$$

The type-specific conditional choice probability for the observed choice y_i is

$$\mathbb{P}_{iljt} = \frac{\exp\left(\widetilde{V}_l[i][y_i]\right)}{\sum_{j=1}^{|\mathcal{J}|} \exp\left(\widetilde{V}_l[i][j]\right)}.$$

Taking logs gives the type-specific log-probability of the observed choice:

$$\log\left(\mathbb{P}_{iljt}\right) = \widetilde{V}_l[i][y_i] - \log\left(\sum_{j=1}^{|\mathcal{J}|} \exp\left(\widetilde{V}_l[i][j]\right)\right).$$

For numerical stability, I compute the logsumexp term using the standard identity: letting $m_{il} = \max_{j \in \mathcal{J}} \widetilde{V}_l[i][j]$,

$$\log\left(\sum_{j=1}^{|\mathcal{J}|} \exp\left(\widetilde{V}_l[i][j]\right)\right) = m_{il} + \log\left(\sum_{j=1}^{|\mathcal{J}|} \exp\left(\widetilde{V}_l[i][j] - m_{il}\right)\right).$$

Subtracting m_{il} before exponentiation prevents overflow when decision utilities are large.

The type-specific panel log-likelihood for household i under type l is the sum of the per-period log-probabilities over all T_i periods in which household i is observed:

$$\ell_{il} = \sum_{t=1}^{T_i} \log\left(\mathbb{P}_{iljt}\right).$$

Since the household's type is unobserved, the mixture log-likelihood for household i integrates over both types using the household-specific mixing weights π_{il}:

$$\log\left(\mathcal{L}_i\right) = \log\left[\pi_{i1} \cdot \exp\left(\ell_{i1}\right) + \pi_{i2} \cdot \exp\left(\ell_{i2}\right)\right].$$

For numerical stability, I compute this using the log-sum-exp identity: letting $m_i = \max\left(\log\left(\pi_{i1}\right) + \ell_{i1},\ \log\left(\pi_{i2}\right) + \ell_{i2}\right)$,

$$\log\left(\mathcal{L}_i\right) = m_i + \log\left[\exp\left(\log\left(\pi_{i1}\right) + \ell_{i1} - m_i\right) + \exp\left(\log\left(\pi_{i2}\right) + \ell_{i2} - m_i\right)\right].$$

The sample log-likelihood is

$$\mathscr{L}\left(\boldsymbol{\theta}^u\right) = \sum_{i=1}^{N} \log\left[\sum_{l=1}^{2} \pi_{il} \cdot \prod_{t=1}^{T_i} \mathbb{P}_{iljt}\right].$$

The optimizer minimizes $-\mathscr{L}\left(\boldsymbol{\theta}^u\right)$. The transition-density term does not appear in this second-stage criterion because the estimation conditions on first-stage transition objects and observed state paths. Moreover, the initial conditions-density term does not appear

either because the estimation conditions on observing the initial addiction state using the method described in Table X.4.6..

X.4.11. *Parameter Optimization*

I estimate the structural utility parameters $\boldsymbol{\theta}^u$ by minimizing the negative log-likelihood using a modified Nelder-Mead algorithm. Throughout this subsection, I write $\boldsymbol{\theta}$ to denote $\boldsymbol{\theta}^u$ for brevity. Nelder-Mead is a derivative-free simplex method, which is well-suited to this setting because the objective function involves solving a value function via iteration at each parameter evaluation, making analytical gradients unavailable and numerical gradients expensive.

The standard Nelder-Mead algorithm maintains a simplex of $(D+1)$ vertices in D-dimensional parameter space, where each vertex $\boldsymbol{\theta}^{(i)} \in \mathbb{R}^D$ has an associated objective value $f^{(i)}$. The first vertex in the simplex is simply the parameter vector under consideration, while the next D vertices slightly modify each respective dimension of the parameter vector to form a $(D+1)$-dimensional simplex. At each iteration, the algorithm attempts to replace the worst vertex (highest objective value) with a better point through a sequence of geometric operations. Let $\boldsymbol{\theta}^{(w)}$ denote the worst vertex, $\bar{\boldsymbol{\theta}}$ denote the centroid of all vertices excluding the worst, and $\boldsymbol{\theta}^{(b)}$ denote the best vertex. By best/worst, I simply mean the vertex corresponding to the best/worst objective value $f^{(i)}$. The operations are:

- *Reflection*: Compute $\boldsymbol{\theta}_r = \bar{\boldsymbol{\theta}} + \left(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^{(w)}\right)$, which mirrors the worst point through the centroid. This calls Nelder-Mead a single time to evaluate $\boldsymbol{\theta}_r$.
- *Expansion*: If the reflected point is the best in the simplex, try $\boldsymbol{\theta}_e = \bar{\boldsymbol{\theta}} + 2(\boldsymbol{\theta}_r - \bar{\boldsymbol{\theta}})$, extending further past the reflected point. This calls Nelder-Mead a single time to evaluate $\boldsymbol{\theta}_e$.
- *Contraction*: If the reflected point is worse than the second-worst vertex (i.e., the reflected point gives an objective value that is the new worst or new second-worst), try $\boldsymbol{\theta}_c = \bar{\boldsymbol{\theta}} + \frac{1}{2}\left(\boldsymbol{\theta}_r - \bar{\boldsymbol{\theta}}\right)$, pulling back toward the centroid. This calls Nelder-Mead a single time to evaluate $\boldsymbol{\theta}_c$.
- *Shrink*: If contraction fails to improve, shrink the entire simplex toward the best vertex by moving each vertex halfway toward $\boldsymbol{\theta}^{(b)}$. This calls Nelder-Mead a D-times to evaluate the remaining D vertices that are not $\boldsymbol{\theta}^{(b)}$.

Notably, if $\boldsymbol{\theta}_r$ is worse than the best vertex $\boldsymbol{\theta}^{(b)}$, but better than the second worst vertex, we simply accept that point by replacing $\boldsymbol{\theta}^{(w)}$ with $\boldsymbol{\theta}_r$ and move on, resulting in a single call to the objective function to evaluate $\boldsymbol{\theta}^r$. Given we are considering a D-dimensional parameter, the minimum number of objective calls by Nelder-Mead is 1, while the maximum is $D+2$, where the D calls come having to do the shrink operation,

and the remaining calls come from a reflection and an expansion or contraction, where the latter two are mutually exclusive.

The algorithm evaluates the objective function at each candidate point and accepts the operation that yields the greatest improvement. Over successive iterations, the simplex contracts around a local minimum. Since Nelder-Mead can become trapped in local minima and is sensitive to the initial simplex, I embed it within a multi-start framework described below.

Let $\boldsymbol{\theta}_0$ denote the vector of starting parameter values and let $\boldsymbol{v} \in \mathbb{R}^D$ denote a vector of initial simplex deviations, where each element v_d controls the size of the initial perturbation along parameter dimension d and D is the dimension of $\boldsymbol{\theta}$. The algorithm consists of $L = 20$ outer tries, each containing $M = 5$ inner tries.

1. For each outer try $l = 1, \ldots, L$:
   (a) Reset the simplex deviations to $\boldsymbol{v}$ at the start of each outer try.
   (b) For each inner try $m = 1, \ldots, M$:
      i. Construct a $(D+1)$-vertex simplex around the current parameter vector $\boldsymbol{\theta}$ using the current simplex deviations.
      ii. Run Nelder-Mead for up to 250 iterations with function value tolerance 10^{-2}.
      iii. Track the best $\boldsymbol{\theta}$ across inner tries.
      iv. Decay the simplex deviations exponentially: $\boldsymbol{v}_m = \boldsymbol{v} \cdot \rho^m \cdot \boldsymbol{u}_m$, where ρ is chosen so that $\rho^{M-1} = 0.10$ and $\boldsymbol{u}_m$ is a vector of independent $\mathbb{U}[0.5, 2.0]$ draws. This shrinks the simplex from full scale at $m = 1$ to 10% of the original deviations at $m = M$.
      v. Perturb the starting point for the next inner try: $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_{\min} + 0.25 \cdot \boldsymbol{v}_m \odot (2\boldsymbol{w} - \mathbf{1})$, where $\boldsymbol{w} \sim \mathbb{U}[0,1]^D$ and $\boldsymbol{\theta}_{\min}$ is the current minimizer. Because $\boldsymbol{v}_m$ shrinks across inner tries, early perturbations explore broadly while later ones fine-tune near the minimum.
   (c) Run Nelder-Mead from the best inner-try parameters for up to 1,500 iterations with function value tolerance 10^{-2}, allowing the algorithm to fully converge from the neighborhood found by the inner loop.
   (d) If the converged objective value is lower than the global best, update the global best parameters $\boldsymbol{\theta}^*$ and objective value.
   (e) Reinitialize the starting point for the next outer try using ϵ-greedy selection with time-decayed exploration. The exploration probability decays exponentially:

$$\epsilon(l) = \epsilon_{\text{end}} + (\epsilon_{\text{start}} - \epsilon_{\text{end}}) \cdot \exp(-\lambda \cdot (l-1)),$$

with $\epsilon_{\text{start}} = 0.75$, $\epsilon_{\text{end}} = 0.25$, and $\lambda = 0.077$.[70] Draw $u \sim \mathbb{U}[0,1]$ and set:

- With probability $\epsilon(l)$: perturb the starting parameters $\boldsymbol{\theta}_0 + \boldsymbol{v} \odot (2\boldsymbol{w} - \mathbf{1})$ (explore).
- With probability $(1 - \epsilon(l))/2$: perturb the global best $\boldsymbol{\theta}^* + \boldsymbol{v} \odot (2\boldsymbol{w} - \mathbf{1})$ (exploit best).
- With probability $(1 - \epsilon(l))/2$: perturb the current minimizer $\boldsymbol{\theta}_{\min} + \boldsymbol{v} \odot (2\boldsymbol{w} - \mathbf{1})$ (exploit current).

All three branches perturb by $\pm\boldsymbol{v}$ so that no restart uses the exact same parameter vector.

The multi-start design addresses the key weaknesses of the Nelder-Mead algorithm: sensitivity to starting values and susceptibility to local minima. The inner loop performs cheap, short explorations with exponentially decaying simplex scale, so early inner tries search broadly while later ones refine locally. The full convergence run then exploits the most promising region found. The ϵ-greedy reinitialization balances exploration (perturbed starting values, favored early) with exploitation (perturbed best or current solution, favored later). The initial simplex deviation vector $\boldsymbol{v}$ is set proportional to the expected magnitude of each parameter to ensure the simplex is well-scaled across dimensions. In practice, the global best objective value stabilizes by approximately $L = 5$ outer tries, with the remaining tries confirming that no better basin exists.

X.4.12. *Hessian Computation*

Standard errors are recovered from the Hessian of the negative log-likelihood at the optimum. Because each evaluation of the objective function requires solving VFI from scratch, analytical gradients are unavailable and automatic differentiation through the VFI fixed-point loop is impractical. I therefore approximate the Hessian entirely via central finite differences on the objective function itself.

Let $f(\boldsymbol{\theta}^u) = -\mathscr{L}(\boldsymbol{\theta}^u)$ denote the negative log-likelihood, $\widehat{\boldsymbol{\theta}^u}$ the estimated parameter vector with D elements, and $\boldsymbol{e}_m$ the m-th standard basis vector. Each parameter receives its own step size

$$h_m = \max\left(\left|\widehat{\theta}^u_m\right| \cdot 10^{-3},\ 10^{-4}\right),$$

[70] At $l = 1$, $\epsilon = 0.75$: 75% of restarts explore by perturbing the starting parameters, while 12.5% exploit the best solution and 12.5% exploit the current solution. By $l = 10$, $\epsilon \approx 0.50$: exploration and exploitation are balanced. By $l = 19$ (the last reinitialization, since $l = 20$ is the final outer try), $\epsilon \approx 0.27$: the algorithm predominantly refines the best solution found so far. The decay rate $\lambda = 0.077$ is chosen so that $\epsilon(10) \approx 0.50$, placing the crossover from exploration-dominant to exploitation-dominant near the midpoint of the outer tries. The high initial exploration rate ensures that the algorithm searches diverse regions of the parameter space early on, while the shift toward exploitation in later tries concentrates effort on refining the most promising basin.

which scales with the magnitude of the estimate (a relative step of 10^{-3}) with a floor of 10^{-4} for near-zero parameters.[71] To reduce numerical noise in the finite differences, the VFI convergence tolerance is tightened to 10^{-6} for the standard error computation (compared to 10^{-4} during estimation). The Hessian $H \in \mathbb{R}^{D\times D}$ is approximated entry by entry as follows.

For diagonal entries:

$$H_{mm} \approx \frac{f\left(\widehat{\boldsymbol{\theta}^u} + h_m \boldsymbol{e}_m\right) - 2f\left(\widehat{\boldsymbol{\theta}^u}\right) + f\left(\widehat{\boldsymbol{\theta}^u} - h_m \boldsymbol{e}_m\right)}{h_m^2},$$

which requires two objective evaluations per parameter (plus the center evaluation $f(\widehat{\boldsymbol{\theta}^u})$, which is computed once).

For off-diagonal entries ($m \neq m'$):

$$H_{mm'} \approx \frac{f\left(\widehat{\boldsymbol{\theta}^u} + h_m \boldsymbol{e}_m + h_{m'} \boldsymbol{e}_{m'}\right) - f\left(\widehat{\boldsymbol{\theta}^u} + h_m \boldsymbol{e}_m - h_{m'} \boldsymbol{e}_{m'}\right) - f\left(\widehat{\boldsymbol{\theta}^u} - h_m \boldsymbol{e}_m + h_{m'} \boldsymbol{e}_{m'}\right) + f\left(\widehat{\boldsymbol{\theta}^u} - h_m \boldsymbol{e}_m - h_{m'} \boldsymbol{e}_{m'}\right)}{4h_m h_{m'}},$$

which requires four objective evaluations per pair. The Hessian is symmetric, so only the upper triangle is computed: $H_{m'm} = H_{mm'}$.

The total number of objective evaluations is $1+2D+4\binom{D}{2}$. For $D = 21$ parameters, this amounts to $1 + 42 + 840 = 883$ evaluations, each requiring a full VFI solve for both types.

The procedure is:

1. Evaluate $f\left(\widehat{\boldsymbol{\theta}^u}\right)$ once to obtain the center value.
2. For each diagonal pair (m, m): evaluate $f\left(\widehat{\boldsymbol{\theta}^u} + h_m \boldsymbol{e}_m\right)$ and $f\left(\widehat{\boldsymbol{\theta}^u} - h_m \boldsymbol{e}_m\right)$, compute H_{mm}.
3. For each off-diagonal pair (m, m') with $m < m'$: evaluate f at the four corner perturbations, compute $H_{mm'}$, and set $H_{m'm} = H_{mm'}$.
4. Invert the Hessian to obtain the variance-covariance matrix:

$$\widehat{\mathbb{V}}\left(\widehat{\boldsymbol{\theta}^u}\right) = H^{-1}.$$

5. Compute standard errors as the square roots of the diagonal:

$$\widehat{\text{se}}\left(\widehat{\theta_m^u}\right) = \sqrt{H_{mm}^{-1}}.$$

[71] A good finite difference step should be proportional to the parameter's magnitude so that the perturbation produces a meaningful change in the objective function. For a large parameter (e.g., $\widehat{\theta}_m^u = 5.0$), the step is $h_m = 0.005$. For a near-zero parameter (e.g., $\widehat{\theta}_m^u = 0.0001$), the proportional step 10^{-7} would be so small that the resulting change in the objective is dominated by numerical noise from VFI convergence. The floor of 10^{-4} prevents this by ensuring the step is always large enough. This is a standard practice in numerical differentiation; see, e.g., Section 5.7 of Press et al. (2007), who recommend scaling the step with the parameter magnitude for the same reason.

The estimation maximizes the likelihood with the flavored habit stock normalized to $[0, 1]$ via $\widetilde{a}^{\text{flav}} = \psi_3 \cdot a^{\text{flav}}$, so the estimated coefficients $\hat{\gamma}_2$, $\hat{\gamma}_3$, and $\hat{\gamma}_4$ are in units of utils per unit of normalized stock. To express them in utils per unit of raw stock a^{flav}, each must be rescaled: $\tilde{\gamma}_k = \hat{\gamma}_k \cdot \hat{\psi}_3$ for $k \in \{2, 3, 4\}$. Since both $\hat{\gamma}_k$ and $\hat{\psi}_3$ are estimated jointly and therefore correlated, the delta method gives

$$\widehat{\text{Var}}(\tilde{\gamma}_k) \approx \hat{\psi}_3^2 \cdot \widehat{\text{Var}}(\hat{\gamma}_k) + \hat{\gamma}_k^2 \cdot \widehat{\text{Var}}\left(\hat{\psi}_3\right) + 2\,\hat{\psi}_3\,\hat{\gamma}_k \cdot \widehat{\text{Cov}}\left(\hat{\gamma}_k,\, \hat{\psi}_3\right),$$

where all variances and covariances are read from the relevant entries of H^{-1}.

The logit parameters $\hat{\pi}_0$ and $\hat{\pi}_{\text{TYA}}$ are estimated directly. The table reports the implied Type 1 probability evaluated at the two extremes of TYA composition, $\bar{h} \in \{0, 1\}$:

$$p(\bar{h}) = 1 - \sigma\big(\hat{\pi}_0 + \hat{\pi}_{\text{TYA}} \cdot \bar{h}\big).$$

The gradient of $p(\bar{h})$ with respect to $(\pi_0,\, \pi_{\text{TYA}})$ is

$$\nabla p(\bar{h}) = -\sigma\big(\hat{\pi}_0 + \hat{\pi}_{\text{TYA}} \cdot \bar{h}\big)\left[1 - \sigma\big(\hat{\pi}_0 + \hat{\pi}_{\text{TYA}} \cdot \bar{h}\big)\right] \begin{pmatrix} 1 \\ \bar{h} \end{pmatrix},$$

so the delta method variance is

$$\widehat{\text{Var}}\big(p(\bar{h})\big) \approx \nabla p(\bar{h})^\top \widehat{\Sigma}_\pi \, \nabla p(\bar{h}),$$

where $\widehat{\Sigma}_\pi$ is the 2×2 submatrix of H^{-1} corresponding to $(\pi_0,\, \pi_{\text{TYA}})$.

X.5. *Counterfactuals*

In this section, I outline the process used to carry out the counterfactual simulations. The only altered components of the above computational details are the flow utilities and value function iteration process. Furthermore, such counterfactual evaluations are done only after obtaining the estimated $\widehat{\boldsymbol{\theta}^u}$ using the estimation approach described previously.

X.5.1. *Present-Bias*

A central question in the addiction policy literature is whether consumers are time-consistent (Becker and Murphy (1988)) or present-biased (Gruber and Köszegi (2001)). The answer matters for policy evaluation: a present-biased consumer overweights the immediate pleasure of consuming a flavored e-cigarette relative to the future cost of deepened addiction, so removing the temptation through a flavor ban may be welfare-improving from the consumer's own long-run perspective, even though it reduces utility in the moment. An exponential discounter, by contrast, has already internalized the

future costs, so a flavor ban unambiguously reduces welfare by constraining the choice set.

Because the model is estimated under exponential discounting ($\beta = 1$), I evaluate each counterfactual under both $\beta = 1$ and $\beta = 0.7$, holding the estimated $\widehat{\boldsymbol{\theta}^u}$ fixed. For $\beta < 1$, the consumer's problem becomes an intrapersonal game between a current self who overweights the present and a sequence of future selves who do the same (O'Donoghue and Rabin (1999)). I assume sophistication: the current self correctly anticipates the future selves' present-biased behavior and optimizes accordingly.[72] This generates two distinct value objects. The decision utility governs what the consumer actually chooses:

$$\overline{v}_{dljt}\left(\boldsymbol{x}_{it}\right) = \overline{u}_{iljt}\left(\boldsymbol{x}_{it}; \widehat{\boldsymbol{\theta}^u}\right) + \beta\delta \int_{\mathcal{X}} V_l\left(\boldsymbol{x}_{i,t+1}\right) \, dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid j, \boldsymbol{x}_{it}; \widehat{\boldsymbol{\theta}^f}\right),$$

where $\beta < 1$ downweights the continuation value, making the consumer act more myopically than the exponential benchmark. The experienced utility governs how the consumer actually evaluates the outcome after the fact:

$$\overline{v}_{eljt}\left(\boldsymbol{x}_{it}\right) = \overline{u}_{iljt}\left(\boldsymbol{x}_{it}; \widehat{\boldsymbol{\theta}^u}\right) + \delta \int_{\mathcal{X}} V_l\left(\boldsymbol{x}_{i,t+1}\right) \, dF_{\boldsymbol{x}}\left(\boldsymbol{x}_{i,t+1} \mid j, \boldsymbol{x}_{it}; \widehat{\boldsymbol{\theta}^f}\right).$$

The difference is the wedge $\beta\delta$ versus δ on the continuation value. The consumer chooses using $\overline{v}_{dl}$ but experiences payoffs according to $\overline{v}_{el}$. When $\beta = 1$, the two coincide and the model collapses to the standard exponential case used in estimation. The sophisticated continuation value V_l solves a fixed point that accounts for this wedge:[73]

$$V_l\left(\boldsymbol{x}_{it}\right) = \sum_{j=1}^{|\mathcal{J}|} \mathbb{P}_{iljt}\left(\boldsymbol{x}_{it}\right) \cdot \overline{v}_{eljt}\left(\boldsymbol{x}_{it}\right) - \sum_{j=1}^{|\mathcal{J}|} \mathbb{P}_{iljt}\left(\boldsymbol{x}_{it}\right) \cdot \log \mathbb{P}_{iljt}\left(\boldsymbol{x}_{it}\right),$$

where the choice probabilities $\mathbb{P}_{iljt}$ are the softmax of decision utility $\overline{v}_{dl}$. The intuition is that the continuation value reflects what the consumer will actually experience (through $\overline{v}_{el}$), but the choices generating those experiences are made by a present-biased self (through $\overline{v}_{dl}$ in the probabilities). This captures the sophisticated consumer's understanding that her future self will overconsume relative to what she would prefer from today's long-run perspective.

The welfare implications of a flavor ban therefore depend on which utility concept is

[72] Under the alternative naive assumption, the consumer incorrectly believes her future selves will behave as exponential discounters, so the continuation value is computed as if $\beta = 1$ will govern future choices even though $\beta < 1$ governs the current choice. Because naive consumers do not anticipate their own future overconsumption, they underestimate the addiction cost of consuming today and therefore consume more than sophisticated consumers. In the context of a flavor ban, naive consumers would exhibit larger pre-ban overconsumption and thus larger welfare gains from the ban. Sophistication is the more conservative assumption for evaluating flavor bans: if the ban is welfare-improving even for sophisticated consumers who partially self-correct, it would be at least as beneficial for naive consumers who do not.

[73] Full derivations and the assumptions required for this representation are provided in Section X.3..

used and on the degree of present bias. Under decision utility, the ban always reduces welfare by constraining the choice set. Under experienced utility, the ban can be welfare-improving if it prevents overconsumption that the consumer's long-run self would prefer to avoid. By evaluating each counterfactual at multiple β values, I trace out how the welfare assessment transitions from unambiguously negative (at $\beta = 1$) to potentially positive (at low β), quantifying the degree of present bias required for the flavor ban to pass a welfare test from the consumer's own long-run perspective.

X.5.2. *Comprehensive Flavor Ban*

To simulate a comprehensive ban on all flavored e-cigarettes, I remove all flavored alternatives from the choice set by setting their flow utilities to $-\infty$. Define the set of banned alternatives as $\mathcal{J}_{\text{ban}} = \{j \in \mathcal{J} : k(j) \in \{E_{f,na},\, E_{f,a},\, B_{f,na},\, B_{f,a}\}\}$, where $E_{f,na}$ and $E_{f,a}$ denote non-FDA and FDA-authorized flavored e-cigarettes, and $B_{f,na}$ and $B_{f,a}$ denote the corresponding flavored bundle categories. For each latent type $l \in \{1, 2\}$, I copy the pre-computed flow utility array and set

$$\overline{u}_l^{\text{ban}}[h, j, a^f, a^s, p] = \begin{cases} -\infty & \text{if } j \in \mathcal{J}_{\text{ban}}, \\ \overline{u}_l[h, j, a^f, a^s, p] & \text{otherwise.} \end{cases}$$

Because $\exp(-\infty) = 0$, banned alternatives receive zero choice probability under the softmax without requiring any modification to the VFI or likelihood routines.

The counterfactual proceeds as follows:

1. Hold fixed the estimated parameters $\widehat{\boldsymbol{\theta}^u}$ and first-stage objects $\widehat{\boldsymbol{\theta}^f}$.

2. For each $\beta \in \{0.7, 1.0\}$:

   (a) Solve the sophisticated VFI (Section X.4.9.) twice in parallel: once using the status quo flow utilities $\overline{u}_l$ and once using the banned flow utilities $\overline{u}_l^{\text{ban}}$, for both types $l \in \{1, 2\}$. This yields four value function arrays: V_l^{sq} and V_l^{ban} for $l \in \{1, 2\}$.

   (b) Compute household-specific posterior type probabilities using the status quo choice probabilities and observed choices:

$$\widehat{\pi}_{il} = \frac{\pi_{il} \cdot \prod_{t=1}^{T_i} \mathbb{P}^{\text{sq}}_{il j_{it} t}(\boldsymbol{x}_{it})}{\sum_{l'=1}^{2} \pi_{il'} \cdot \prod_{t=1}^{T_i} \mathbb{P}^{\text{sq}}_{il' j_{it} t}(\boldsymbol{x}_{it})},$$

   where π_{il} is the prior mixing weight. The posterior is computed from the status quo, not the counterfactual, to ensure type assignment reflects observed behavior.

(c) Forward-simulate both scenarios for $T_{\text{sim}} = 36$ months using $S = 100$ Monte Carlo draws per household. For each household i, draw d, and period t:

i. Assign type: draw $l \sim \text{Bernoulli}\left(\widehat{\pi}_{i2}\right)$ using a pre-drawn uniform u_i^d.

ii. Interpolate $\overline{v}_{dl}^{\text{sq}}$ and $\overline{v}_{dl}^{\text{ban}}$ at the household's continuous state $\left(h_i, a_{it}^f, a_{it}^s, p_{iCt}, p_{iEt}\right)$ using quadrilinear interpolation over the $\left(a^f, a^s, p_C, p_E\right)$ grid.

iii. Compute softmax choice probabilities for each scenario:

$$\mathbb{P}_j^{\text{sq}} = \frac{\exp\left(\overline{v}_{dlj}^{\text{sq}}\right)}{\sum_{j'} \exp\left(\overline{v}_{dlj'}^{\text{sq}}\right)}, \qquad \mathbb{P}_j^{\text{ban}} = \frac{\exp\left(\overline{v}_{dlj}^{\text{ban}}\right)}{\sum_{j'} \exp\left(\overline{v}_{dlj'}^{\text{ban}}\right)}.$$

iv. Sample choices using the same pre-drawn uniform u_{it}^d via inverse CDF for both scenarios.

v. Update addiction stocks deterministically based on each scenario's sampled choice:

$$a_{i,t+1}^{f,\text{sq}} = (1 - \psi_2) \cdot a_{it}^{f,\text{sq}} + n_{j^{\text{sq}}},$$
$$a_{i,t+1}^{f,\text{ban}} = (1 - \psi_2) \cdot a_{it}^{f,\text{ban}} + n_{j^{\text{ban}}},$$

and analogously for the slow stock. Prices evolve via the AR(1) process using the same pre-drawn normal shocks for both scenarios.

(d) Aggregate simulated choices into per-period category shares, mean addiction $\bar{a}_t = \left(\bar{a}_t^f + \bar{a}_t^s\right)/2$, and mean welfare, separately for TYA-present and TYA-absent households.

The use of common random numbers ensures that differences between the status quo and ban scenarios reflect the policy change rather than simulation noise. No structural parameters are re-estimated; the counterfactual change is entirely through the restriction of the choice set.

X.5.3. *FDA-Authorized Flavor Ban*

The FDA-authorized flavor ban follows the identical procedure as the comprehensive ban, except the set of banned alternatives is restricted to FDA-authorized flavored products only:

$$\mathcal{J}_{\text{ban}}^{\text{FDA}} = \{j \in \mathcal{J} : k(j) \in \{E_{f,a},\, B_{f,a}\}\}.$$

Non-FDA flavored alternatives ($E_{f,na}$ and $B_{f,na}$) remain in the choice set with their original flow utilities. All other steps (VFI, posterior computation, forward simulation, and aggregation) are identical.

X.5.4. *Flavor Tax*

The flavor tax counterfactual applies a per-milliliter excise tax $\tau > 0$ to all flavored e-cigarette and flavored bundle alternatives. Unlike the ban, which sets flow utilities to $-\infty$, the tax shifts flow utilities downward by an amount proportional to the taxed quantity. The tax enters as an increase in effective expenditure. Because $\omega_E < 0$, the shift is negative, making flavored alternatives less attractive.

The tax shift is applied uniformly across all state points (h, a^f, a^s, p) because it does not depend on the household's addiction level, TYA status, or current prices: it is a fixed per-unit charge on the flavored quantity. All remaining steps (VFI, posterior computation, forward simulation, aggregation) are identical to the ban counterfactual, with $\overline{u}_l^{\text{tax}}$ replacing $\overline{u}_l^{\text{ban}}$ in the VFI. Because no alternatives are removed from the choice set, all $|\mathcal{J}|$ alternatives retain positive choice probability under the tax.

X.6. *Addiction Decay Parameters*

The two-component addiction stock requires calibrating two decay parameters: ψ_1 for the slow (behavioral) component and ψ_2 for the fast (physiological) component. I set $\psi_1 = 0.10$ and $\psi_2 = 0.90$, so that the fast stock loses 90% of its value each month while the slow stock loses only 10%. Both values are grounded in the biomedical literature on nicotine addiction.

When a smoker quits, the brain undergoes a measurable chemical adjustment. Chronic nicotine exposure causes the brain to produce extra nicotinic receptors, which are the proteins that bind to nicotine and generate its rewarding effects. These additional receptors create a neurochemical demand for nicotine that produces the familiar withdrawal symptoms of irritability, anxiety, and difficulty concentrating when not met. Brain imaging studies have tracked how quickly these extra receptors disappear after a smoker quits. Mamede et al. (2007) used such imaging to measure receptor levels at several points after the quit date: at ten days of abstinence, levels were still roughly 26% above those of nonsmokers, but by 21 days they had returned to nonsmoker levels. A larger study by Cosgrove et al. (2009) confirmed the general pattern, but found that normalization can take up to 6 – 12 weeks in some brain regions, with levels still approximately 15% elevated at four weeks. The clinical experience of withdrawal mirrors this receptor timeline, with symptoms typically peaking within the first week and largely resolving within three to four weeks McLaughlin, Dani, and De Biasi (2015). Setting $\psi_2 = 0.90$ means that the fast stock retains only 10% of its value after one month and is effectively zero after two months, matching the receptor evidence that the acute neurochemical disruption of quitting is largely over within one to two months.

If receptor levels return to normal within a few weeks, why do so many ex-smokers relapse months or even years later? The answer lies in a second, slower process: the learned associations between everyday cues and the act of smoking. Over years of

use, the brain forms strong connections between nicotine reward and the situations in which smoking occurs such as a morning coffee, a work break, stress, social settings. These associations are encoded in the brain's reward and memory systems and persist long after acute chemical dependence has resolved. Hughes, Keely, and Naud (2004) document that the relapse curve among untreated quitters follows a hyperbolic shape: most relapses occur in the first days and weeks, but the hazard rate remains meaningfully elevated for months and years. Among smokers who achieve one full year of abstinence, which is well past the point where receptor levels have normalized, roughly 10% relapse in the following year Hughes, Peters, and Naud (2008). Krall, Garvey, and Garcia (2002) find that even among those who remain abstinent for two or more years, 2 – 4% relapse annually through years two to six, with the rate not falling below 1% until after a decade of abstinence.

Setting $\psi_1 = 0.10$ yields a slow stock that retains 53% of its value at six months, 28% at one year, and 8% at two years. This trajectory is consistent with the relapse evidence: the stock is large enough in the first year to generate the elevated relapse rates documented by Hughes, Peters, and Naud (2008), declines gradually through years two to six in line with the Krall, Garvey, and Garcia (2002) estimates, and approaches zero only after several years. A faster decay rate such as $\psi_1 = 0.40$ would drive the slow stock to near zero within six months, leaving the model unable to account for the substantial fraction of relapses that occur after that point.

X.7. *How Do Packs of Cigarettes Translate to Milliliters of E-Liquid*

It is difficult to understand how a single pack of cigarettes translates to a single mL of e-liquid. However, one can get a sense for how consumption is interpretable across units through the amount of nicotine consumed and absorbed. For instance, although a typical cigarette may contain 10 - 14 mg of nicotine, empirical evidence suggests that only about 1 - 1.5 mg is absorbed systemically per cigarette (Benowitz, Hukkanen, and Jacob (2009)), roughly only 10% of total nicotine concentration. Thus, I normalize nicotine consumption through cigarettes to 12 mg per cigarette and I normalize nicotine absorption to 1.25 mg per cigarette. The former implies one pack of cigarettes contains 240 mg of nicotine while the latter implies one pack of cigarettes is equivalent to $20 \times 1.25 = 25$ mg of nicotine absorbed. Using this cigarette-equivalent standard, I define consumption across alternative products, e-cigarettes and cessation goods, in relation to nicotine absorption through cigarettes.

E-cigarettes commonly contain about 50 mg/mL nicotine, and with roughly 50% absorption (Yingst et al. (2019); Prochaska, Vogel, and Benowitz (2022)). This implies approximately 25 mg absorbed per mL. Let d denote the nicotine concentration in mg/mL and e_d denote the volume of e-liquid with dosage d consumed in mL. The total absorbed nicotine from e_d mL of e-liquid with a dosage of d nicotine concentration in mg/mL is $0.5 \times d \times e_d$ mg. I assume each cigarette delivers approximately 1.25 mg of

absorbed nicotine as discussed above, so the total absorbed nicotine from n cigarettes is $1.25n$ mg. Equating the absorbed nicotine from e-liquid and cigarettes and solving for n gives

$$0.5d \times e_d = 1.25n$$
$$n = 0.4d \times e_d.$$

Thus, e_d mL of e-liquid with nicotine concentration dosage d mg/mL is equivalent to $n = 0.4 \times d \times e_d$ cigarettes in terms of absorbed nicotine. For example, if $d = 50$ mg/mL and $e_d = 1$ mL, then $n = 0.4 \times 50 \times 1 = 20$ implying that consuming 1 mL of e-liquid at 50 mg/mL delivers the same absorbed nicotine as approximately 20 cigarettes, or a single pack of cigarettes. I use this conversion to translate consumption of mL of e-liquid into cigarette-pack equivalents.

It is important to emphasize that cigarette-pack equivalence throughout is defined exclusively in terms of absorbed nicotine, not total nicotine content. Because combustible cigarettes deliver nicotine inefficiently—absorbing only a small fraction of nicotine content—one pack of cigarettes (240 mg nicotine content) corresponds to only 25 mg of absorbed nicotine. By contrast, e-cigarettes deliver nicotine more efficiently. As a result, an e-cigarette quantity that delivers the same absorbed nicotine as one cigarette pack need not contain the same total nicotine content. For example, under the assumptions above, consuming 1 mL of e-liquid at 50 mg/mL delivers 25 mg of absorbed nicotine—equivalent to one cigarette pack in absorbed terms—even though it contains only 50 mg of nicotine content, or roughly one-fifth of the nicotine content in a cigarette pack. Throughout the analysis, all quantities expressed in cigarette or pack equivalents refer to this absorbed-nicotine scale.

Table A.8. Nicotine Consumption vs. Absorption Equivalence for E-Cigarettes

E-cigarette use	Nicotine content (packs)	Nicotine absorbed (packs)
1 mL at 50 mg/mL	$50/240 \approx 0.21$	$25/25 = 1$
4.8 mL at 50 mg/mL	$240/240 = 1$	$120/25 = 4.8$

X.8. *Why Reduced-Form Evidence is Insufficient*

The reduced-form evidence above is informative, but not sufficient for the policy questions this paper addresses. As documented in Section II.2., flavor restrictions have spread rapidly across states and localities, and understanding the effects of a national flavor ban requires projecting beyond any variation observed in the data. The three key counterfactuals in this paper, a comprehensive flavor ban, a ban limited to FDA-authorized brands, and a per-milliliter excise tax on flavored e-cigarettes, are either hypothetical at the federal level or have never been implemented in any form, so their effects are inherently out of sample and cannot be recovered from reduced-form methods

alone.

Beyond the counterfactual problem, the state-level bans that do exist were enacted largely contemporaneously with the COVID-19 pandemic, which disrupted tobacco purchasing patterns in ways that are difficult to separate from the regulatory treatment. This confounding makes it hard to draw clean causal inferences from the existing natural experiments. A structural approach avoids this problem by estimating preference and addiction parameters, then simulating policy effects using the estimated model rather than relying on post-ban variation in outcomes.

The structural model also enables analysis that reduced-form methods cannot support by design. Nicotine addiction creates dynamic linkages between current choices and future utility, so a policy that reduces consumption today lowers future addiction stocks, which in turn reduces the cost of continued abstinence, compounding the initial effect over time. Tracing these long-run addiction trajectories requires a model that explicitly tracks the evolution of the addiction state. Similarly, a flavored habit stock means that flavor bans do not simply reduce demand uniformly: they disproportionately affect consumers who have built up flavored consumption habits, and the speed at which those habits decay determines how persistent the policy's effects are. Neither of these dynamics is recoverable from reduced-form estimates of short-run substitution. Finally, how present-biased preferences interact with addiction and flavor habits reshapes both the behavioral response to bans and the welfare interpretation of the results: under present bias, part of the measured welfare cost of a ban reflects an internality correction rather than genuine consumer harm, a distinction that requires the full structural framework to quantify.

A dynamic structural model overcomes these limitations. It computes both short- and long-run effects by solving the consumer's full dynamic problem and decomposes behavioral mechanisms that are separately identified through the model's structure.